



Computer Networking Essentials

An essential guide to understanding networking
theory, implementation, and interoperability

MICROSOFT CORP.
EXHIBIT 1005



Computer Networking Essentials

Debra Littlejohn Shinder

Cisco Press

Cisco Press
201 West 103rd Street
Indianapolis, IN 46290 USA

Computer Networking Essentials

Debra Littlejohn Shinder

Copyright © 2002 Cisco Systems, Inc.

Published by:

Cisco Press

201 West 103rd Street

Indianapolis, IN 46290 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America 3 4 5 6 7 8 9 0

Third Printing January 2002

Library of Congress Cataloging-in-Publication Number: 2001090429

ISBN: 1-58713-038-6

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

This book is designed to provide information about basic networking and operating system technologies. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The author, Cisco Press, and Cisco Systems, Inc., shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers’ feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book or otherwise alter it to better suit your needs, you can contact us through e-mail at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher
Editor-in-Chief
Executive Editor
Cisco Systems Management

Production Manager
Development Editor
Senior Editor
Copy Editor
Technical Editor
Reviewers

Associate Editor
Cover Designer
Composition
Indexer

John Wait
 John Kane
 Carl Lindholm
 Michael Hakkert
 Tom Geitner
 William Warren
 Patrick Kanouse
 Kitty Wilson Jarrett
 Jennifer Chisholm
 Jill Batistick
 Dr. Thomas W. Shinder
 Lynn Bloomer
 Wayne Jarvimaki
 Michael R. Hanson
 Shannon Gross
 Louisa Klucznik
 Steve Gifford
 Tim Wright



Corporate Headquarters
 Cisco Systems, Inc.
 170 West Tasman Drive
 San Jose, CA 95134-1706
 USA
<http://www.cisco.com>
 Tel: 408 526-4000
 800 553-NETS (6387)
 Fax: 408 526-4100

European Headquarters
 Cisco Systems Europe
 11 Rue Camille Desmoulins
 92782 Issy-les-Moulineaux
 Cedex 9
 France
<http://www-europe.cisco.com>
 Tel: 33 1 58 04 60 00
 Fax: 33 1 58 04 61 00

Americas Headquarters
 Cisco Systems, Inc.
 170 West Tasman Drive
 San Jose, CA 95134-1706
 USA
<http://www.cisco.com>
 Tel: 408 526-7660
 Fax: 408 527-0883

Asia Pacific Headquarters
 Cisco Systems Australia, Pty.,
 Ltd
 Level 17, 99 Walker Street
 North Sydney
 NSW 2059 Australia
<http://www.cisco.com>
 Tel: +61 2 8448 7100
 Fax: +61 2 9957 4350

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2001, Cisco Systems, Inc. All rights reserved. Access Registrar, AccessPath, Are You Ready, ATM Director, Browse with Me, CCDA, CCDE, CCDP, CCIE, CCNA, CCNP, CCS1, CD-PAC, *CiscoLink*, the Cisco NetWorks logo, the Cisco Powered Network logo, Cisco Systems Networking Academy, Fast Step, FireRunner, Follow Me Browsing, FormShare, GigaStack, IGX, Intelligence in the Optical Core, Internet Quotient, IP/VC, iQ Breakthrough, iQ Expertise, iQ FastTrack, iQuick Study, iQ Readiness Scorecard, The iQ Logo, Kernel Proxy, MGX, Natural Network Viewer, Network Registrar, the Networkers logo, *Packet*, PIX, Point and Click Internetworking, Policy Builder, RateMUX, ReyMaster, ReyView, ScriptShare, Secure Script, Shop with Me, SlideCast, SMARTnet, SVX, TrafficDirector, TransPath, VlanDirector, Voice LAN, Wavelength Router, Workgroup Director, and Workgroup Stack are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Empowering the Internet Generation, are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, Cisco, the Cisco Certified Internetwork Expert Logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Collision Free, Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastLink, FastPAD, IOS, IP/TV, IPX, LightStream, LightSwitch, MICA, NetRanger, Post-Routing, Pre-Routing, Registrar, StrataView Plus, Stratum, SwitchProbe, TeleRouter, are registered trademarks of Cisco Systems, Inc. or its affiliates in the U.S. and certain other countries.

All other brands, names, or trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0010R)

Introduction

Computer Networking Essentials helps you understand the fundamentals of computer networking concepts and implementation and introduces you to the client and server operating systems that run on networked PCs.

Concepts covered in this book include the history of networking, networking terminology, networking theory and established standards, and implementation of local-area and wide-area networks. Special emphasis is placed on understanding network protocols and how they operate at all layers of the networking model. Emphasis also is placed on the interoperability of networks that run on multiple protocols, platforms, and operating systems.

Specialty areas such as security, remote access, virtual private networking, thin client networking, monitoring, management, and troubleshooting are covered thoroughly. Emerging technologies that are expected to impact the future of networking are also introduced.

Who Should Read This Book

This book's primary audience is professionals who are beginning training in the networking industry and those who need a review of basic concepts.

The secondary audience includes corporate training faculties and staff and members of the business world who work with information technology personnel and require a broad overview of the concepts involved in networking from the small business to the enterprise-level corporation.

A third target audience is the general user who wants to know more about how computers communicate over networks. The book's approach is designed to be user-friendly and accessible to the non-technical reader who is overwhelmed by the jargon found in vendor documentation and technical manuals.

This Book's Organization

This book is organized into four parts and includes 19 chapters, an appendix, and a glossary. The following sections describe the contents of each part of the book.

Part I: Introduction to Networking Concepts

Chapter 1, "Introduction to PC Networking," introduces you to the basic concepts of PC networking by providing a brief history of electronic communications and networking and a summary of where PC networking is today.

Chapter 2, "Categorizing Networks," discusses the categorization of networks according to physical scope, administrative model, network operating system, protocols in use, topology, and architecture.

Chapter 3, "Networking Concepts, Models, and Standards," provides an overview of binary communications and introduces two popular networking models: the Department of Defense (DoD) model on which the TCP/IP protocols are based and the Open Systems Interconnection (OSI) model, which was developed by the International Organization for Standardization (ISO). Specifications set forth by the Institute of Electrical and Electronics Engineers (IEEE) and vendor-specific models are also covered.

Chapter 4, "Networking Communications Methods," discusses signaling methods and provides an understanding of analog, digital, broadband, baseband, asynchronous, synchronous, simplex, duplex, and multiplexed signaling. Media access methods are described, including CSMA/CD, CSMA/CA, token passing, and demand priority.

Chapter 5, "LAN Links," discusses popular LAN types, including Ethernet, Token Ring, FDDI, AppleTalk, and ARCnet.

Chapter 6, "WAN Links," provides an overview of WAN connections such as PSTN, ISDN, t-carriers, Frame Relay, X.25, and CATV network, as well as high-speed connectivity solutions such as ATM, SONET, and SMDS. This

chapter also covers LAN-to-WAN connection solutions, including Internet Connection Sharing (ICS), Network Address Translation (NAT), proxy servers, and routed connections.

Part II: Networking Hardware and Software

Chapter 7, “Physical Components of the Network,” introduces students to the many types of networking media, including coax, twisted-pair cable, and fiber-optic cable, as well as to wireless technologies such as laser, infrared, radio, and satellite/microwave communications. Connectivity devices such as repeaters, hubs, bridges, routers, and switches are also discussed.

Chapter 8, “Networking Protocols and Services,” describes common LAN protocols—TCP/IP, NetBEUI, IPX/SPX—and discusses the OSI protocol suite. PPP and SLIP, which are WAN link protocols, and PPTP and L2TP, which are common tunneling protocols, are also presented.

Chapter 9, “The Widest Area Network: The Global Internet,” discusses the evolution of the Internet, the protocols used for Internet communications—HTTP, FTP, NNTP, SMTP, and POP—and the TCP/IP protocol suite.

Chapter 10, “Network Operating Systems,” discusses general network administration practices and then looks at the specifics of common server operating systems, including Windows NT, Windows 2000, NetWare, UNIX, and Linux.

Chapter 11, “Directory Services,” describes the Directory Services Protocol (DAP) and the Lightweight Directory Access Protocol (LDAP), as well as the X.500 standards developed by the ISO to promote directory services compatibility and interoperability. Novell’s NDS, Microsoft’s Active Directory, and Banyan VINES’ StreetTalk directory services are covered in some depth.

Chapter 12, “Desktop Operating Systems,” looks at the client side of the client/server network and discusses the advantages and disadvantages of common desktop clients, such as DOS, Windows, Linux, Macintosh, and OS/2, and how each can be integrated into popular NOS environments.

Chapter 13, “Hybrid Networks,” provides information about interoperability solutions and protocol gateways that allow PCs running different operating systems, protocols, and platforms to communicate with one another. This chapter also looks at PC-to-mainframe communications using Systems Network Architecture (SNA) solutions.

Part III: Network Specialty Areas

Chapter 14, “Protecting the Network,” addresses security issues and provides an overview of basic cryptography concepts, public and private key encryption, certificate services, firewalls and proxies, and internal security measures such as “smart cards” and advanced authentication technologies. It also provides guidance for developing security policies for your network. The second half of the chapter discusses disaster recovery plans, including implementation of disk fault tolerance (or RAID), regular scheduled backups, and server clustering.

Chapter 15, “Remote Access,” discusses methods of connecting to a server from a remote location using remote connectivity devices such as modems, ISDN terminal adapters, and customer premises equipment (CPE) for dedicated lines. Dial-in server configuration and special security considerations are also covered.

Chapter 16, “Virtual Private Networking,” provides an overview of VPN concepts and discusses the tunneling protocols used to provide VPN security.

Chapter 17, “Thin Client Networking,” discusses Network Computers, Net PCs, and Windows-based terminals. Windows terminal services, Citrix Metaframe, web-based computing, the X Window system and Java virtual machines—and the role each plays in thin client networking—are also discussed.

Chapter 18, “Monitoring, Management, and Troubleshooting Tools,” presents an introduction to the TCP/IP utilities and other tools built into the various operating systems. This chapter also examines commercial products such as Sniffer Pro, LANalyzer, Microsoft’s Systems Management Server, Novell’s ManageWise, and IBM’s Tivoli.

Part IV: The Future of Networking

Chapter 19, “Tomorrow’s Technologies,” takes a look into the future of PC networking. It discusses ways of overcoming the current limits of IP, including the new version of IP—IPv6. The goal of universal connectivity is addressed, and more exotic possibilities such as artificial intelligence, quantum computing, and cybernetic life forms are presented as possible components of tomorrow’s networks.

This Book’s Features

This book contains several elements that help you learn about operating systems and networking:

- **Figures, listings, and tables**—This book contains figures, listings, and tables that help to explain concepts, commands, and procedural sequences. Diagrams illustrate network layouts and processes, and screenshots assist students in visualization configuration procedures. In addition, listings and tables provide summaries and comparisons of features and characteristics.
- **Author’s notes, tips, sidebars, and cautions**—These elements are included to provide you with extra information on a subject. You will probably find these asides to be very beneficial in real-world implementations.
- **Chapter summaries**—At the end of each chapter is a summary of the concepts covered in the chapter, which provides a synopsis of the chapter and can serve as a study aid.
- **Further Reading**—Each chapter includes a list of resources for additional information about the topics covered in the chapter, including website URLs and books and articles that cover the topic in more detail.
- **Review questions**—After the Further Reading section in each chapter are 10 review questions that serve as an end-of-chapter assessment. The questions are designed to reinforce the concepts introduced in the chapter and to help students evaluate their understanding before moving on to the next chapter.

The conventions used to present command syntax in this book are the same conventions used in the *Cisco IOS Command Reference*, as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In examples (not syntax), boldface indicates user input (for example, a **show** command).
- *Italics* indicates arguments for which you supply values.
- Square brackets [] indicate optional elements.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Braces and vertical bars within square brackets—for example, [x {y | z}]—indicate a required choice within an optional element. You do not need to enter what is in the brackets, but if you do, you have some required choices in the braces.

WAN Links

The technologies, media, and equipment that work well for the short distances spanned by a LAN or MAN are generally not suitable for long-distance wide-area networks (WANs). In today's very mobile world, high-performance, cost-effective WAN technologies are a necessity for many reasons:

- Executives and other employees need access to their corporate networks while on the road or at home.
- Companies with branch offices in widely dispersed geographic locations need network connectivity between locations.
- Organizations want to share information with other organizations physically separated by long distances.
- Commercial, governmental, and educational bodies and individuals need access to the resources available on the global Internet.

It is obviously impossible to string Ethernet cable from the home office in Denver to the branch office in Houston. Even if cabling distance limitations did not apply, this would not be a viable solution for connecting international sites.

WANs require a whole new set of technologies and rules of implementation. In this chapter, we discuss the concept of networking over long distances and the technologies commonly used to connect computers that are located in different states, countries, or even different continents. These range from the Public Switched Telephone Network (PSTN) already in place in most of the world to modern high-tech solutions such as satellite communications technologies that enable us to “talk” to computers in space.

Wide-area networking presents many challenges not encountered in implementing a network that is confined to one geographic area. A WAN is *not* just a really big LAN. Rather, it is a collection of many separate LANs, connected by links that are different in many ways from LAN links. WANs that span international boundaries require consideration of even more factors, including time zones and language differences.

Designing a WAN is a complex task. Choosing the appropriate technology involves analyzing the purpose(s) the WAN will serve, the number of users, the bandwidth requirements, and the patterns of use. We can categorize these considerations as follows:

- WAN hardware
- WAN topologies
- Network switching types
- New and emerging WAN technologies
- LAN/WAN connectivity

We look at each issue in the sections that follow.

WAN Hardware

The hardware necessary to implement a WAN link can be as simple and inexpensive as a telephone line and a modem at each end. On the other hand, it can be complex and costly. In general, equipment cost and complexity increases with increased speed and reliability.

In the following sections, we discuss common WAN devices, including modems, ISDN and digital subscriber line (DSL) terminal adapters, and customer premises equipment (CPE) used with dedicated links such as T-carrier connections and X.25.

Modems

To establish a network connection (to an Internet service provider or to a dial-up server on a private network) over public telephone lines, you use a device called a *modem*.


TIP The word “modem” is derived from the actions it performs; a modem *modulates* and *demodulates* a signal. In other words, it converts the sending computer’s digital signal to analog for transmission over the analog line and then converts it back to digital for processing by the receiving computer.

Modems come in two physical types: *internal* and *external*. Each has advantages and disadvantages, and configuration is slightly different depending on the type. Either way, modems are *serial* devices, which means bits are sent one at a time. This can be contrasted with *parallel* devices, such as printers, to which multiple bits can be sent simultaneously. A serial transmission is analogous to a group of people marching in a straight line, and a parallel transmission is like having the same group marching in rows of three across.

Internal Modems

One advantage of the internal modem is compactness. It is a circuit board card that fits in an ISA (Industry Standard Architecture) or PCI (Peripheral Component Interconnect) slot inside the computer, as shown in Figure 6-1. This means that you don't have to find room for an extra device on your desk. In addition, you are not required to buy a serial cable, which you might be forced to do if you use an external model that doesn't include one in the box.

Figure 6-1 *An internal modem is a circuit board that fits inside the computer.*



Internal Modem Configuration Parameters

Internal modems are traditionally more difficult to configure than external modems. You must set the IRQ, the input/output addresses, and the virtual com ports to ensure that they don't conflict with the settings of some other device in your computer. Let's look more closely at each setting and how it is used:

- **Interrupt Request (IRQ)**—This is an assigned location that designates where the system expects the device to interrupt it when the device sends a signal. Signals from different devices that go to the processor on the same interrupt line would interfere with each other, so a separate IRQ must be assigned to each device.
- **Input/Output (I/O) address**—This is the location where data sent from the device is stored before it is processed by the CPU. As with the IRQ, if multiple devices attempt to use the same I/O address, one or both devices might not work properly.
- **Virtual com port**—This is a logical port number, by which the operating system identifies a serial port. You must set each serial device to use a different com port.

All popular operating systems provide a means by which you can view how resources are being used, and which ones are not in use, so that you can choose free resources to assign to your new device.

Changing the Internal Modem Settings

Internal modems generally provide a way to change the configuration settings. Depending on the manufacturer and model, you can change IRQ, I/O, and com port settings with the following:

- **Dip switches**—These are small switches on the circuit board that can be moved to a different position. The position of the switch designates which setting is to be used.
- **Jumpers**—Pairs of metal pins built into the circuit board, these represent an electrical contact point. Jumpers are configured by placing a small plug on the pins to complete the circuit. The instructions that come with your internal modem tell you how the jumpers should be set to use a specific IRQ, I/O address, or com port.
- **Software**—Some modems do not have physical switches or jumpers, but do come with a software program that is run to change the configuration.

Plug and Play

Many modern modems support *Plug and Play (PnP)* technology, which enables the operating system to detect the device, install the necessary software drivers, detect what resources are free on the computer, and assign those resources to the device automatically. Little or no intervention is required from the user.

PnP is great—when it works and when you are aware of a few caveats. If you buy a modem or other device that is advertised as Plug and Play, it is automatically configured *only* if the following is true:

- Your computer's BIOS (Basic Input/Output System) supports PnP.
- You are running a PnP operating system.

Both criteria must be met. Computer motherboards produced after 1995 usually support PnP. Operating systems that support PnP include Windows 95, 98, ME, and 2000.

NOTE


Windows NT is *not* a PnP operating system. However, it does have limited PnP functionality and detects some modem types.

External Modems

External modems have a couple advantages over the internal variety:

- Most external modems provide status lights, which indicate when the modem is powered on, connected, or transferring data. See Figure 6-2.
- External modems are generally easier to install and configure. There are no switches or jumpers to set, and you don't have to open the computer case.

Figure 6-2 *External modems provide status indicator lights.*



External modems require power cords to plug into an electrical outlet, but internal modems run off the computer's power. A serial cable connects the modem to one of the serial ports on the back of the computer.

Serial Port Considerations

To use an external modem, you need a free *serial port*. Most computers have two built-in serial ports, labeled COM 1 and COM 2, with connectors on the back of the computer.

Many devices, such as scanners, digital cameras, and serial pointing devices, also use serial ports. If your computer does not have a free serial port, you have a few options:

- Use an internal modem.
- Install an expansion card in your computer, which enables you to add serial port connections.
- If your computer's motherboard supports universal serial bus (USB), you can chain multiple serial devices, such as modems, off a single serial port. You might have to add a card to provide a USB connector, and you need a USB modem.

UART Chips

Serial ports use a chip called a UART (Universal Asynchronous Receiver/Transmitter) to handle serial communications. This chip comes in different types, and the type used determines how fast data can be transferred over that serial port.

The first PCs had 8250 UART chips. The top speed for this chip is 9600 bps, which means that even if you attach a high speed (56 kbps) modem to one of these ports, your speed would be limited by the UART.

Modern computers have UART chips in the 16450 or 16550 series. These serial ports can support transfer speeds of up to 115,200 bps.

16650 and 16750 UART chips are also available as add-on “enhanced serial port” cards. Internal modems have their own UART chips built into the card, so the speed of the computer’s com port is irrelevant.

NOTE If you have a high-speed modem and a modern computer, but are able to connect only at low speeds, check the com port configuration settings. Some operating systems set the com ports to 9600 bps by default; you need to change this setting to realize the port’s full capacity.

Modem Drivers

Drivers are software programs that act as a liaison between the hardware device and the operating system. Driver software is usually supplied by the modem manufacturer with the device, or it can be downloaded from the manufacturer’s Web site.

You must install the correct driver software for your device because if operating system code included support for all hardware devices that could possibly be used with it, the operating system would require significantly more disk space—much of it wasted on driver software that would never be used.

Modem Configuration

In addition to installing the driver software that enables the operating system to recognize the modem, and setting the IRQ, I/O address, and com port that the modem will use, you have to configure the modem to dial and maintain a connection. Modern operating systems have built-in support for dialup networking. You might have to install the remote access services if the modem was not present when the operating system was installed.

Modem Banks

A computer can be configured as a *dialup server* (also called a *remote access server*) to enable other computers to dial into it and connect to it over the phone lines. Computers running powerful server software can support many incoming remote access connections simultaneously; for instance, Windows NT Server supports up to 256 connections.

How can you connect 256 modems to a remote server? When you have many simultaneous dial-in connections (for example, when the server belongs to a company with many telecommuters who need to connect to the corporate network from home), you can use a *modem bank*. Modem banks are also called *modem nests* or *modem pools*.

A modem bank enables you to use a group of modems (usually mounted together in a rack) with a single server, and host multiple remote connections. The rack of modem cards is controlled by an interface that connects to the server, to a router, or directly to the local network. Of course, you need a phone line for each separate connection.

ISDN and DSL Adapters

The device used to connect a computer to an Integrated Services Digital Network (ISDN) or DSL telephone line is often referred to as a modem. It is more accurately called a *terminal adapter* because it does not modulate and demodulate signals because ISDN lines are digital, unlike the analog PSTN lines.

ISDN Adapters

ISDN adapters, such as modems, come in both internal and external varieties. They are configured similarly to modems, but the typical 128-kbps ISDN service consists of two data channels that each run at 64 kbps. The two channels are commonly used in a *multilink* configuration to provide the 128-kbps bandwidth. We discuss ISDN technology later in this chapter in the section, “ISDN.”

The two data channels have separate telephone numbers in most cases. ISDN adapters are configured with information about the *service profile identifier (SPID)* for each channel, which consists of the telephone number, a two-digit sharing terminal identifier, and a two-digit terminal identifier (TID). Some modern ISDN adapters support automated SPID selection and do not require you to enter this information.

DSL Adapters

Both ends of a DSL connection require a device called an *endpoint* (and often referred to as a *DSL modem*), which connects to an Ethernet NIC installed in the computer. In some cases, the endpoint/modem is external. In others, the endpoint and the NIC are placed together on the same card.

Customer Premises Equipment

Customer premises equipment (CPE) is a general term that encompasses several different devices. The customer’s site requires this hardware to process incoming transmissions from WAN links such as T-carrier lines, X.25 connections, and Frame Relay links.

Common types of CPE include the following:

- A channel service unit/digital service unit (CSU/DSU), used with circuit-switched connections such as a T-1 line. The CSU receives and transmits signals to and from the WAN line. The DSU manages line control, timing errors, and signal regeneration.
- A packet assembler/disassembler (PAD), used with packet-switched connections such as X.25. The PAD is an asynchronous device that enables multiple terminals to share a network line. Users dial into PADs through modems.

WAN Topologies


We discussed LAN topologies in Chapter 2, “Categorizing Networks,” and some of those same concepts apply to WANs. In the context of WANs, however, the *topology* describes the arrangement of the transmission facilities.

The simplest WAN topology is a simple point-to-point connection. The WAN, like the LAN, also can use traditional networking topologies such as a ring or star.

The Point-to-Point WAN

A point-to-point WAN is similar to the LAN topology referred to as a linear bus. A remote access link, which can be anything from a 56-kbps dial-up modem connection to a dedicated T-1 line, connects each point on the WAN to the next. See Figure 6-3 for an illustration of this.

Figure 6-3 A point-to-point WAN directly connects two endpoints.



This is a relatively inexpensive way to connect a small number of WAN sites. However, it is not fault tolerant. For example, in Figure 6-3, if the equipment at the Dallas office fails, San Francisco and Boston cannot communicate with one another. Limited scalability (the capability to “grow gracefully,” that is, to continue to function efficiently as the network grows larger) is another disadvantage. If you add another point to the WAN at Nashville, between Dallas and Boston, you increase the number of hops required for Boston to communicate with Dallas or San Francisco.


NOTE In wide-area networking, a *hop* is defined as the trip from one router to the next. The *hop count* is the number of routers the packet passes through from source to destination.

The point-to-point link works best for small WANs with only two or three locations.

The WAN Ring

A ring is constructed by establishing a point-to-point connection from Point A to Point B, from Point B to Point C, and from Point C back to Point A, as shown in Figure 6-4.

Figure 6-4 A WAN can use a ring topology.



The ring topology provides redundancy. In the example shown in Figure 6-4, if the line between Dallas and San Francisco goes down, data can still be transferred between the two cities by going through Boston.


The ring topology is more expensive to implement than the single point-to-point topology, and it suffers from the same scalability problem as the point-to-point topology.

The ring topology works well for WANs that connect only a few locations and that need the reliability offered by the redundant pathways.

The WAN Star

When a WAN is laid out in a star configuration, a device called a *concentrator router* is used; it serves as a central point to which all network routers are connected. In Figure 6-5, for example, the concentrator router is located at the Dallas headquarters.

Figure 6-5 A WAN arranged in a star topology is scalable.



The star topology is more scalable than the ring, and in a star, it is easier to add locations to the WAN.

The disadvantage of the star is its single point of failure. In the case shown in Figure 6-5, this point of failure is the Dallas concentrator router. If this device fails, communications cease among all points on the network.

Full- and Partial-Mesh WANs

A mesh topology, where there are multiple connections between points, provides the most fault-tolerant and reliable WAN. Unfortunately, it is also the most expensive to implement, and it becomes cumbersome if there are a large number of sites.


A full mesh topology requires that every site in the network be connected to every other site. With a partial mesh, a smaller number of redundant connections exist. This provides reliability approaching that of the full mesh, at significantly lower cost.

Refer to Figures 2-10 and 2-11 in Chapter 2 for an illustration of mesh and partial (hybrid) mesh topologies.

Multitiered WANs

A multitiered WAN is similar to the star in that it uses concentrator routers, but is more reliable because it links two or more of these concentrators with other locations “cascaded” off the main routers. See Figure 6-6 for an illustration.

Figure 6-6 A multitiered WAN offers more reliability than a simple star.



The multitiered WAN is scalable because additional locations (and even additional tiers) can be added to the network relatively easily. This topology is used for large, fast-growing networks.

Traffic flow can become a problem on large, multitiered WANs. Flow patterns should be carefully analyzed to ensure the most efficient placement of equipment for best performance.

Network Switching Types

Data can travel to a distant destination over several different types of lines, using one of two popular switching technologies:

- Circuit switching
- Packet switching

In the next sections, we examine how each technology works.

Circuit Switching Versus Packet Switching

Many people confuse circuit switching and packet switching. An example of a circuit-switched network is the telephone network. When you place a call to a branch office in Boston, a circuit is established for the duration of the call. The signal uses that circuit, or pathway, until you break the connection by hanging up. If you call again tomorrow, a different pathway might be taken.

An example of a packet-switched network is the Internet. When you send an e-mail message to the Boston office, it is broken down into chunks called *packets*. Each packet might take a different pathway to reach the destination computer; they are reassembled at the other end.

Circuit-Switching Networks

The first switching type we examine is the *circuit-switched* network.

Dialup Versus Dedicated Connections

Circuit-switched networks include both dialup and dedicated leased lines.

A dialup connection is a *temporary* connection, established for the duration of the session. Dialup connections can, however, be implemented as “always-on” connections with which you dial up the remote server and then do not end the connection. With a dialup connection, it is possible to hang up and dial a different location if you choose to do so. For example, you can end the connection to one ISP and then dial in and connect to another.

A dedicated connection is one that goes only from one specific point to another (for example, from your business office to your ISP).

Circuit-switching technology has a long history and is older than packet switching. Circuit switching is most appropriate when data must be transmitted in real time (as with a

telephone conversation). Circuit-switched networks are connection-oriented networks because a connection is established before transmission begins.

We briefly examine the following circuit-switched technologies in the following subsections of this chapter:

- PSTN
- ISDN
- DSL
- Leased lines
- Digital data service (DDS)
- T-carriers
- Switched 56

PSTN

The most common type of dialup WAN link is made by using the PSTN—the analog phone lines that are installed in most residences and businesses.

PSTN has two big advantages:

- It is available almost all over the world.
- It is inexpensive.

A dialup connection using ordinary phone lines is easy to implement. Besides a modem, no special equipment is required, and analog modems are readily available, simple to configure, and inexpensive.

The telephone system was not originally created with data transfer in mind; it was designed to transmit voice. High speed was not an issue, so there is an inherent limit to the attainable transfer rate.

Line quality is also a factor. Even with top-of-the-line 56-kbps modems, many telephone lines are capable of providing no more than 40 kbps–45 kbps.

ISDN

ISDN was designed to eventually replace POTS and provide a reliable digital connection suitable for both voice and data. Although that hasn't happened, and the recent advent of faster technologies at lower cost means it probably won't, ISDN still offers some advantages.

The characteristics of ISDN include the following:

- As its name implies, ISDN is a digital link. Because it does not have to convert data from digital to analog format and back, performance and reliability are high.
- It is more readily available than some of its newer competitors, such as DSL.
- Although it is a dialup technology, ISDN can be used as an always-on link (that is, dedicated ISDN).
- ISDN service is more expensive than analog service (PSTN) and requires specialized equipment, both at the telephone company central office (that is, at the digital switch) and at the customer's premise (that is, at the ISDN terminal adapter).

An ISDN circuit is made up of one or more channels that carry data (called bearer channels, or B channels) and a control channel (called the Delta channel, or D channel).

Each B channel provides 64 kbps of bandwidth, and B channels can be aggregated by using *inverse multiplexing*. This enables you to combine the bandwidth of multiple channels to create one high-speed connection. The D channel provides either 16 or 64 kbps, depending on the interface implementation.

ISDN is offered by most telephone companies in two standard access interfaces:

- **Basic Rate ISDN (BRI)**—This interface consists of two 64-kbps B channels (for an aggregate usable bandwidth of 128 kbps) and one 16-kbps D channel.
- **Primary Rate ISDN (PRI)**—This consists of 23 64-kbps B channels (for an aggregate bandwidth of 1.472 Mbps) and one 64-kbps D channel.

BRI is often implemented for residential or small business high-speed data transfer, and PRI is commonly used for digital voice transmission in conjunction with private branch exchange (PBX) telephone systems.

NOTE

A *PBX* is a private telephone network operated within an organization. Internal users share outside lines, and calling within the organization requires dialing only a four-digit extension. A traditional PBX required a switchboard operator, who answered all incoming calls and then routed each to the appropriate extension. Modern equipment automates this process.

DSL

DSL is a relatively new technology, offered by telephone companies as an add-on service over existing copper wires. DSL offers several advantages over other WAN link types.

The following list contains characteristics of DSL:

- DSL offers speeds up to and exceeding those of T-1, at a fraction of the cost. In many areas, DSL service costs less than ISDN.
- DSL is an always-on technology. There is no need to dial up each time you wish to connect.
- Both voice and data can be transmitted over the same line simultaneously.
- At present, availability is limited. The telephone company central office (CO) that is servicing the location must have DSL equipment installed, and for most “flavors” of DSL, you must be within a specified number of feet from the CO to get DSL service.

DSL comes in several varieties:

- **ADSL (Asymmetric DSL)**—This is the most common implementation. Speeds vary from 384 kbps to 6 Mbps (or more) downstream, typically combined with a lower upstream speed.
- **SDSL (Symmetric DSL)**—This provides the same speed for downloads and uploads.
- **HDSL (High Data Rate DSL)**—This variety typically provides bandwidth of 768 kbps in both directions.
- **VDSL (Very High Data Rate DSL)**—This is capable of bandwidths between 13 Mbps and 52 Mbps.
- **IDSL**—This is DSL over ISDN lines. It has a top speed of 144 kbps, but is available in areas that don’t qualify for other DSL implementations.

The generic term for DSL, encompassing all implementations, is *xDSL*.

ADSL, currently the most popular DSL implementation, generally provides a fast downstream transfer rate (typically 1.5 Mbps) and a slower upstream rate. This is based on the theory that most users of the Internet primarily access e-mail and surf the Web, which are download-intensive tasks. The lower upload rates do not work as well, however, if you wish to host a Web or FTP server, or engage in other upload-intensive tasks.

ADSL typically uses frequency-division multiplexing (FDM) to split the bandwidth and create multiple channels. Some ADSL implementations use a different method, called *echo cancellation*. It is more efficient, but also more complex and more costly.

Table 6-1 summarizes currently available DSL implementations.

Table 6-1 *Comparison of DSL Implementations*

DSL Type	Average Speeds	Advantages	Disadvantages
ADSL	384 kbps to 6 Mbps (downstream)	Relatively inexpensive; more widely implemented than other types.	Can be installed only within 17,500 ft of a telephone company CO; upstream speed is usually much slower.
SDSL	Up to 3 Mbps	Offers the same data rate upstream and downstream.	Generally more expensive and less widely available than ADSL.
IDSL	144 kbps	Can be installed in many locations where other DSL types are not available because of distance.	More expensive than ADSL; considerably slower speed.
HDSL	768 kbps up and downstream	Faster than IDSL and some implementations of ADSL.	Not widely available.
VDSL	13 Mbps to 52 Mbps	Extremely high speed for live audio and video.	Not widely available; most expensive DSL type.

Leased Lines

For WAN links that require guaranteed high performance and reliability, an option is to lease lines from the telephone company for private use. A leased line provides a permanent connection from one point to another (for example, from one branch office to another, or from your company LAN to your ISP).

DDS

DDS was one of the first digital services made available to the public. DDS provided a 56-kbps transfer rate. It lost the popularity contest to T-carrier technology because a T-1 line typically provides more bandwidth per dollar.

T-carriers

T-carriers are dedicated digital circuits that are typically leased by large companies to provide high-speed data, voice, audio, and video over a highly reliable point-to-point connection.

T-carrier circuits are typically established over copper wires, but they can also run over fiber-optic cable, coaxial cable, and even wireless technologies.

A CSU/DSU is used at each end of the connection to encode the data to be sent over the T-carrier.

Although prices for T-1 lines have fallen dramatically over the last decade, it is still an expensive option. A T-1 line typically costs 10 to 20 times that of DSL service for comparable speed (1.5 Mbps).

Why would anyone pay for T-1 when low-cost high-speed options are available? Availability itself is one reason; DSL has only recently become widespread, and even in areas where it is offered, many businesses and residences are not within the distance limitations required to order the service.

Another reason to pay extra for T-1 is guaranteed bandwidth. This is called the *committed information rate (CIR)*. With 1.5-Mbps DSL service, the telephone company sells you a service that has a maximum transfer rate of 1.5 Mbps. Your line might or might not actually perform at that speed at a given time. When you lease a T-1 connection, the telephone company guarantees the data rate of 1.544 Mbps. This can be an important consideration for corporate enterprises that depend on network performance.

The T in T-carrier refers to transmission channel; the signal itself is more accurately referred to as the data signal (DS) rate. You hear both DS-1 and T-1 used to refer to the same line type. Table 6-2 lists common T-carrier implementations.

Table 6-2 Common T-Carrier Implementations

Carrier Designation	Data Signal Rate	Data Transfer Speed
T-1	DS-1	1.544 Mbps
T-2	DS-2	6.312 Mbps
T-3	DS-3	44.736 Mbps
T-4	DS-4	274.760 Mbps

T-carrier lines consist of multiple 64 kbps channels. It is possible to lease just part of a T-1 line (in 64 kbps increments) if you don't need the entire 1.544 Mbps bandwidth; this is called *fractional T-1*.

Switched 56

Switched 56 is an enhanced version of PSTN. It is a digital switched-circuit connection that transfers data over one 56 kbps channel. Switched 56 is less expensive—but also much slower—than a T-1 line. It is a dialup technology, thus it can be appropriate in cases in which a dedicated connection is unnecessary.

Unlike PSTN, which is theoretically capable of data rates of 56 kbps but in practice generally attains speeds no higher than 50 kbps (and often falls far short of that), Switched 56 can provide a reliable full 56 kbps connection. Because it is a digital connection, error rates are lower than on a regular PSTN (analog) line.

The connection is called “switched” because individual 56 kbps channels are switched out of a T-1 circuit and sent to a specific user location.

The popularity of Switched 56 has suffered as other low-cost, high-bandwidth options (such as ISDN and DSL) have become more widely available.

Packet-Switching Networks

Packet-switching networks are networks in which data packets can take different routes to reach the same destination. At the receiving end, the packets are put back together in the correct order. Packet-switched networks are often depicted as a cloud because the exact route of travel of the data is unknown.

Packet-switching technologies include the following:

- X.25
- Frame Relay
- Asynchronous Transfer Mode (ATM)

NOTE Although you often hear the term *X.25 network*, the technically correct term is *Public Switched Data Network (PSDN)*. X.25 is the protocol that is used for communication between the data terminal equipment and the network.

We discuss each in more detail in the following sections.

X.25

X.25 was one of the first packet-switching networks and was designed to work with IBM mainframes, such as the IBM 360, and use analog transmission.

X.25 was originally called the ARPAnet 1822 protocol; the name X.25 came from the specifications for the protocols used by this technology, established by the International Telegraph and Telephone Consultative Committee (CCITT) in 1976.

NOTE The CCITT changed its name to the International Telecommunications Union (ITU) in 1993.

PSDN technology operates at the first three layers of the OSI model. The technology called X.25 is actually made up of several protocols:

- PSDN uses a protocol called X.21 at the physical layer (a variation of the X.21 physical layer protocol, X.21bis, is used in the United States).
- A protocol named Link Access Procedure Balanced (LAPB) is used at the data link layer.
- At the network layer, the Packet Layer Protocol (PLP) is used to assemble frames from the data link layer into packets.

The primary objective in designing the X.25 protocol was reliability. At the time it was designed in the 1970s, both the computers in use and the telephone lines were prone to error. Thus, the PSDN running on X.25 included redundant error-checking to compensate for these problems. The result was a highly reliable means of data transfer, but performance was slowed by the extra error-checking activity. The PSDN usually transfers at 64 kbps or below.

There are still public switched data networks in use today. To make a WAN connection over a PSDN, you can do one of the following:

- Dial into a packet assembler/disassembler (PAD) with an asynchronous modem
- Make a synchronous connection using the X.32 protocol
- Use an X.25 smart card to connect directly to the PSDN

Frame Relay

Frame Relay is a newer packet switching technology, which was designed to be used over digital lines and which grew out of X.25. Frame relay is a variation on and improvement to the X.25 technology, developed by the CCITT. Frame Relay uses only the first two OSI layers rather than the first three (as X.25 does). It was developed to take advantage of modern computers and telephone lines, which are far more reliable than those in use when X.25 originated. It has become a popular option for WAN links, and it generally offers a higher-performance, more cost-effective solution than does X.25.

Frame Relay operates only at the two lowest levels of the OSI model, the physical and data link layers. Frame Relay uses less overhead than X.25, and thus, it is faster. Frame Relay can run at T-1 and T-3 speeds (from 1.5 Mbps to almost 45 Mbps).

Frame relay is called a *fast packet* technology.

A typical Frame Relay implementation uses a *permanent virtual circuit (PVC)* to provide an always-on connection. Because service providers generally charge fees based on usage (referred to as *bandwidth on demand*), you can avoid the cost of a dedicated leased line.

Frame Relay has high performance because it does not include the extensive error checking and correction of X.25. Frames with errors are discarded, and it is up to the endpoints

(communication computers) to detect the missing packet and request retransmission. Because transmission is digital, there are relatively few errors to contend with, and Frame Relay works well in a WAN environment over T-1 lines.

ATM

ATM is a popular packet-switching technology that was designed to support high-speed applications such as streaming audio and video. An important concept for ATM networks is quality of service (QoS), which is a way to control the allocation of network bandwidth to specific applications to provide guaranteed bandwidth where it is most important.

ATM is hardware based, which means that all equipment on the network must be designed to work with ATM. The advantage is that this results in high speeds for processing and switching. Standard ATM transfer rates are 25 Mbps, 155.520 Mbps, and 622.080 Mbps, and ATM is capable of speeds of 10 Gbps. Unfortunately, that performance comes at a high price. ATM is expensive to implement because all network hardware must support ATM, and network interface cards (NICs), hubs, and other ATM-compatible equipment is costly.

ATM is a modern digital technology that breaks data into 53-byte fixed-length units called *cells*. Five bytes are used for the ATM header, which contains addressing information.

ATM can be used for both LANs and WANs, and it uses multiplexing to transfer voice, data, and video simultaneously over the network. Cell switching is a function of the ATM hardware (unlike Frame Relay, in which it is a software function).

You can connect to an ATM network through a direct connection or an on-demand connection. The connection between the two endpoints is a *virtual circuit*; it can be either a PVC or a switched virtual circuit (SVC). With either a PVC or an SVC, ATM uses predefined circuits instead of establishing the virtual circuits at the time of connection, as X.25 and Frame Relay do. This saves a great deal of time and is another factor in ATM's high speed.

Many networking experts predict that in the future, ATM will become the technology of choice for both LANs and WANs.

Emerging WAN Technologies

New, faster, and more efficient WAN technologies are being developed all the time. Many of these interoperate with one another to provide support for the high-bandwidth applications in use today and those expected to be in demand in the future.

In the following sections, we look at new high-speed technologies, including OC-SONET, Broadband ISDN, CATV, and SMDS.

OC-SONET

OC stands for *optical carrier*, and SONET stands for *Synchronous Optical Network*. SONET is a physical layer protocol that provides for high-speed transmission using fiber-optic media. SONET is capable of rates of almost 20 Gbps, and ATM can run over SONET to achieve very high data transfer speeds.

NOTE You might see the term Synchronous Digital Hierarchy (SDH) used to refer to the SONET technology outside the United States.

The SONET signal rate is measured by OC standards. Table 6-3 illustrates the available transmission rates (called optical carrier levels).

Table 6-3 *OC Signal Transmission Rates*

OC Level	Signal Transmission Rate
OC-1 (base rate)	51.84 Mbps
OC-3	155.52 Mbps
OC-12	622.08 Mbps
OC-48	2.488 Gbps

SONET is used as the physical basis for another technology, broadband ISDN, which is discussed in the next section.

Broadband ISDN

Broadband ISDN (BISDN) is an emerging technology designed to use fiber-optic cable and radio waves to transmit data at high speeds over SONET, FDDI (the Fiber Distributed Data Interface), and Frame Relay.

Broadband technologies, which can send multiple channels of data, video, and voice over the same medium, are growing in popularity as Internet connectivity and other high-bandwidth network usage increases. Other broadband technologies include DSL and cable modem.

CATV

Cable TV (CATV) companies saw a great opportunity: They already had a vast infrastructure of coaxial cable in most major cities and many rural areas, and this cable

could be used not only to transmit television signals, but also to transmit computer data. Numerous cable providers now offer Internet access accounts.

Cable is not a general WAN technology. It was originally designed to enable you to communicate only with the cable company/service provider's server (in the form of receiving incoming television channels). Although customers were all connected to the same network through the coax cable running through their neighborhood, the network was not designed to enable them to communicate with one another. In fact, the network was not designed to enable its users to send data at all—only to receive it. Cable modem changes all that.

Cable Internet access requires a cable modem that connects both to the incoming coax cable and to a NIC in the user's PC (typically this is a 10BaseT Ethernet NIC).

In this scenario, the cable company is also the user's ISP. There is no option to separate the provision of the physical line from the access service as there is with access over telephone lines. In other words, you cannot lease the line from the cable company and use it to connect to some other ISP's server.

On the other hand, when you pay a telco for the use of a phone line (whether an analog PSTN line or a dedicated T-1 line), you can purchase an Internet account from any ISP you choose (including the phone company that provides the line). Although this same method is technically possible with cable, the cable companies have packaged the two services together and the terms of their service contracts require that you use the cable company as your ISP.

Cable infrastructure can support either one-way or two-way transmissions. *One-way cable* provides only downstream transmission over the coax. Uploading must be done over a regular analog phone line that also plugs into the cable modem. With one-way cable, upload speeds are limited to standard rates attainable over PSTN, which is less than 56 kbps. Download speeds vary from 364 kbps to 1.5 Mbps.

Two-way cable provides both uploads and downloads over the coax. Nonetheless, many cable companies limit the upstream speed to 128 kbps to discourage customers from running servers (which is often prohibited by the CATV terms of service contract).

Cable is an always-on technology, but one-way cable still requires you to dial up to establish a connection. A big advantage of CATV is its low cost; however, in many areas, users experience reliability problems. Because cable is a "shared-bandwidth" technology (that is, the entire bandwidth of the cable is divided between all users on that cable segment at a given time), performance might degrade as more users in the neighborhood are added to the network. There are also security issues that, at this time, make CATV more viable for residential use than for business.

SMDS

Switched Multimegabit Data Service (SMDS) is a new packet-switching technology that is designed especially for WAN links that experience a lot of “bursty” traffic. (*Bursty* refers to transmission that comes in “bursts” rather than in a constant, even stream.)

SMDS is connectionless; that is, there is no requirement that a connection or circuit be established before transmitting the data. It uses relatively large packets, up to 7168 bytes in length. SMDS addresses, which are ten-digit numbers (such as a telephone number), are used to identify the SMDS subnetwork. SMDS links are connected to an SMDS switch on the telephone company’s backbone network, typically by multiple OC-3 SONET links.

SMDS was designed as a public network to provide services similar to those of a LAN, except that it spans a metropolitan area. Data transfer speeds typically range from 1.544 Mbps to 45 Mbps. It is scalable and can be used in conjunction with ATM. However, SMDS is not as widely available as Frame Relay and other services, and SMDS equipment may be more difficult to find.

Wireless WAN

In many cases, it is impossible—or at least inconvenient or expensive—to run a wired link to connect WAN sites. Wireless solutions are especially appropriate when it is important that data be communicated in real time, or when users are on the move. Wireless works best for communicating small amounts of data.

The wireless technologies used for WANs include the following:

- **Radio frequency (RF) technologies**—Specialized Mobile Radio (SMR) provides data rates of 1200 bps to 19,200 bps. Enhanced SMR (ESMR) is the digital implementation.
- **Satellite technologies**—This provides both circuit-switched and packet-switched services at speeds of 4800 to 9600 bps.
- **Microwave technologies**—This technology uses cellular techniques over microwave frequencies to provide higher speed and capacity (wireless broadband).
- **Cellular technologies**—This provides a circuit-switched connection over analog or digital cellular links.
- **Packet data network technologies**—This technology provides a packet-switched WAN with no call setup involved.

Compared to wired links, wireless communications are often more costly and relatively slow. For example, analog cellular systems typically provide no more than 14,400 bps transfer rates, while digital cellular offers up to 64 kbps.

LAN/WAN Connectivity

In today's wired world, local connectivity often is not enough. No LAN is an island, or at least, fewer and fewer of them are islands as it becomes vital to business interests that a LAN be able to communicate with the outside. This can mean connecting the LAN to a corporate WAN, the global Internet, or both.

There are several ways to connect your LAN to the outside world, depending upon your budget and needs. Of course, the most obvious way to provide network users with access to other networks is to equip each PC with a modem and phone line. In this manner, each user can establish a dialup connection to an ISP or other remote server. However, this solution has many drawbacks:

- It becomes prohibitively expensive as the number of users increases. Not only must you purchase hardware (the modem) for every computer, you must also pay for a separate telephone line *and* if users are to connect to the Internet, a separate ISP account for each.
- Allowing users to dial out using a modem can create serious security risks if the nature of the data on your network is confidential. The company has little control over which networks the user connects to and the audience to which the company's data might be exposed.
- A high degree of user sophistication is required, which means a significant expense for training users to configure and manage their own dialup connections.

There is a better way. In fact, several alternatives offer advantages over the old-fashioned way of connecting LAN users to a WAN. Each has advantages and disadvantages, and which is best depends on your particular situation.

In the following sections, we briefly discuss the following LAN/WAN connectivity options:

- Translated connections
- Proxy servers
- Routed connections

Translated Connections

One of the most cost-effective ways to connect all computers on a small LAN to the Internet or to another WAN link is through address translation. *Address translation* enables all computers to access the WAN through a single host computer, using only one telephone line and ISP account (or other WAN link) and only one registered public IP address.

How Address Translation Works

A computer running address translation software sits between the public WAN and the private LAN. It has interfaces to both networks. This computer has a private IP address used for communications with other computers on the LAN and a public IP address (which can be assigned through Dynamic Host Configuration Protocol [DHCP] from an ISP's server at the time the WAN connection is established). We refer to this computer as the *address translation host*.

NOTE

DHCP is a service that automatically assigns IP addresses and other TCP/IP settings to computers that are configured to use DHCP. You will learn more about DHCP in Chapter 8, "Networking Protocols and Services."


Address translation works by mapping the private IP address of each computer on the LAN that sends data "outside" to a port number on the host computer. This information is added to the IP header of the packet, which is then sent out over the WAN with the IP address of the host computer (the one that has the physical connection to the WAN) as the source address.

When a computer on the local network opens a Web browser and sends a request to view a URL, for example, the host computer assigns a port number to that request, which identifies the original sending computer. Then the host sends the request out to the ISP's Web server. When the page is returned to the host computer (whose IP address is listed in the header as the source of the request), the host consults its address translation table, matches up the packets with the computer that originally sent the request, and then forwards the Web page to that computer.

The information in the address translation table includes the following:

- The original source and destination IP addresses (identifying the sending computer within the network and the computer outside the network to which the data is sent)
- The original source and destination port numbers (identifying the application making or receiving the request; for example, HTTP requests for Web pages are normally sent to TCP port 80)
- Sequence numbers (identifying the order in which the packets are sent)
- A timestamp

Network address translation (NAT) is the common term for which standards have been developed and published as RFC 1631. Not all address translation technologies comply with these standards. Figure 6-7 illustrates the steps in the address translation process.

Figure 6-7 *The NAT process involves translating private addresses to a public address.*

- 1 The user at the client computer (IP address 192.168.1.9) opens a Web browser application and enters the URL `www.tacteam.net` into the address box. The browser software sends an HTTP request to the IP address associated with the `www.tacteam.net` “friendly name.”
- 2 The NAT host on the client’s LAN maps the request from 192.168.1.9 for `www.tacteam.net` to a port number in the mapping table. The table contains the original source and destination IP addresses and original source and destination TCP/UDP port numbers.
- 3 The NAT host changes the header so that to the outside network, the packet appears to originate not from 192.168.1.9, but from the public IP address assigned to the NAT host’s external network adapter by the ISP.
- 4 The NAT host sends the request for `www.tacteam.net` to the ISP server. Domain Name System (DNS) maps the name to the IP address of the server on which the `www.tacteam.net` homepage is stored.
- 5 The request is received by the `www.tacteam.net` server and the page is returned to the public IP address used by the NAT host.
- 6 The NAT host consults its address translation table to determine whether the page should be sent to the client at 192.168.1.9 (and the TCP/UDP port number to which it should be sent).

NAT Software

Some operating systems, such as Windows 98/2000 and current versions of Linux, have built-in support for address translation. In Linux it is referred to as *IP masquerading*. Windows 98 and 2000 Professional call the feature *Internet Connection Sharing (ICS)*, although it can be used to share a connection to a private remote network or to a VPN as well. Windows 2000 Server supports the more flexible and robust version of address translation, NAT. If an operating system does not include address translation, you can use an add-on NAT program to provide the same functionality. Examples of these programs include:

- **Sygate, from Sybergen Software**—www.sygate.com
- **NAT32, from A.C.T. Software**—www.nat32.com

Another, more sophisticated (and slightly more difficult to configure) type of software that you can use to share a connection is a *proxy*. We discuss this option in the next section.

Proxy Servers

A proxy server does more than provide a shared connection to the WAN, although it does serve this purpose. A proxy server acts as an intermediary, separating the LAN from the outside network, and it can provide protection by filtering incoming and outgoing packets. It also enhances Web performance by *caching* often-requested Web pages.

Proxies use an address translation method, but do not necessarily comply with the NAT specifications in RFC 1631.

How Proxies Work

The proxy server receives requests for Internet resources from proxy clients, similar to the way in which NAT works. The proxy server checks its filter settings (which are configured by the administrator). If the request meets filter requirements, the server looks first in its cache of stored pages. If the requested page is there, the proxy server returns the page to the requesting client. There is no need to send the request on to the ISP server. If the page is not there, the proxy server requests the page from the ISP server, receives it, and returns it to the client.

As with NAT, the internal clients that access the Internet through a proxy server are invisible to the Internet; all outside communication is done by the proxy server.

Proxy Software

Proxy software typically provides more protection and performance enhancement than NAT. However, it might also be more expensive, and it is typically more difficult to set up

because the Internet applications on all the client machines (such as the Web browser) must be individually configured to use a proxy server.

To use NAT, you need only set the client's TCP/IP configuration to obtain an IP address through DHCP, and other necessary information can be automatically distributed to the clients by the DHCP server.

Numerous proxy server applications are available for popular operating systems, and they include the following:

- **Rideway, from DGL**—dgl.com/rideway
- **Winproxy, from Osis**—www.winproxy.com/toc
- **Microsoft Proxy Server, from Microsoft**—www.microsoft.com
- **Squid for UNIX, from SCO**—www.sco.com

Some proxy programs, such as Microsoft Proxy Server, run only on a network operating system (NOS) such as Windows NT or Windows 2000 Server. Others, such as Rideway and Winproxy, can be used on desktop operating systems such as Windows 95/98, Windows NT Workstation, and Windows 2000 Professional.

Software that combines proxy and NAT technologies are sometimes referred to as *transparent proxies*.

Routed Connections

A *routed connection* is another way to provide the computers on the LAN with access to a WAN. A routed connection enables each computer to participate directly on the Internet, unlike NAT and proxy connections, where the individual computers must go through an intermediary.

Configuring a routed connection requires extensive knowledge of TCP/IP addressing, and you must purchase and configure a *router*. Additionally, every computer on the LAN that connects to the outside network must have a “legal” registered public IP address.

NOTE See Chapter 8 for more information on TCP/IP and IP addressing and Chapter 9, “The Widest Area Network: The Global Internet,” for more information on routing.

Why Use a Routed Connection?

A big advantage of the NAT and proxy solutions is the capability to connect the small LAN to a public WAN by using only a single IP address. However, this might not be the best solution in some situations.

Because of the way address translation works, protocols that do not store the addressing information in the IP header do not work with NAT. In some cases, *NAT editors* can be added to make modifications to the IP packet so that NAT will work. In other cases (for example, when packets are authenticated and encrypted using IP Security [IPSec]), address translation is not possible.

Configuring a Routed Connection

A routed connection requires either a dedicated routing device such as a router or the use of a computer running an operating system that enables IP forwarding (in the latter case, the computer acts as the router).

Computers that use TCP/IP to communicate must have the following properties configured:

- An IP address that is valid for the network on which they will communicate
- A subnet mask that designates what part of the IP address identifies the computer and what part identifies the network

Computers participating on a *routed* network must also have a *default gateway* configured. This is the address of the router, which has two network connections: one to the LAN and one to the outside network.

NOTE

The address for a DNS server must be entered if you want to use “friendly” host names (for example, URLs such as `www.tacteam.net`) instead of IP addresses. You learn more about DNS in Chapter 8.

To set up a routed connection to the Internet, the TCP/IP protocol on the router is configured with an IP address, a subnet mask, and a DNS server address obtained from the ISP, and a static default route is configured to use the Internet interface.

The computers on the LAN that connect to the Internet are likewise configured with IP addresses, a subnet mask, and a DNS server address obtained from the ISP. They are also configured to use the IP address of the router on its LAN interface as their default gateway address.

Summary

In this chapter, you have learned about wide-area networking and the established and emerging technologies that can be used to connect computers in distant locations.

We discussed WAN hardware, and then we moved on to WAN topologies and discussed the advantages and disadvantages of WAN configurations such as the point-to-point connection, the ring, the star, the full or partial mesh, and the multitiered WAN.

We then delved into the differences between circuit-switched and packet-switched networks, and we discussed the characteristics and technologies associated with PSTN, ISDN, xDSL, DDS, T-carriers, X.25, Frame Relay, ATM, OC-SONET, BISDN, CATV, SMDS, and wireless WAN technologies.

We next discussed how to connect a LAN to a WAN. Specifically, we learned several ways to provide all computers on a local network with access to the Internet or another outside network. We learned about translated connections that use NAT, which is sometimes called connection sharing or IP masquerading. Then we discussed how to set up a routed connection for those circumstances in which address translation is undesirable or impossible.

This chapter wraps up Part I of this book. You should now have a grasp of basic networking concepts. In Part II, we look at the hardware and software that makes the network run. Chapter 7 introduces you to the components of the physical network.

Further Reading

An excellent resource for information on various WAN technologies is WANsites, at www.networkcomputing.com/wansites/default.html.

Thorough, clear explanations of the different WAN technologies can be found at the High Performance Networking Unleashed Web site, at www.officewizard.com/books/network.

An excellent, detailed discussion of NAT is available online at www.suse.de/~mha/linux-ip-nat/diplom.

Review Questions

The following questions test your knowledge of the material covered in this chapter. Be sure to read each question carefully and select the *best* correct answer or answers.

- 1 Which of the following hardware settings commonly must be configured on an internal modem?(Select all that apply.)
 - a I/O address
 - b IP address
 - c IRQ
 - d Virtual com port

- 2 The CSU/DSU and the PAD are examples of which of the following?
 - a ISDN terminal adapters
 - b Customer premises equipment
 - c Terminal identifiers
 - d Concentrator routers
- 3 PSTN, ISDN, DSL, DDS, and T-carrier links are all examples of what type of network?
 - a Circuit-switched networks
 - b Packet-switched networks
 - c Switched 56 networks
 - d LANs
- 4 Which of the following are true of ISDN? (Select all that apply.)
 - a It is an analog link.
 - b It is generally more expensive than PSTN.
 - c It requires special equipment at the CO and the customer's premise.
 - d It uses a circuit consisting of only one channel.
 - e It can transfer both voice and data.
- 5 Which WAN topology is the most scalable?
 - a Point-to-point.
 - b Ring.
 - c Star.
 - d All of the above are equally scalable.
- 6 What was the first packet switching technology that was based on the ARPAnet 1822 protocol?
 - a Frame Relay
 - b X.25
 - c ATM
 - d DSL

- 7 Which of the following are characteristics of Frame Relay that distinguish it from X.25? (Select all that apply.)
- a Frame Relay offers high performance.
 - b Frame Relay uses packet switching.
 - c Frame Relay uses digital signaling.
 - d Frame Relay does not include extensive error checking.
- 8 Which of the following describes ATM? (Select all that apply.)
- a It uses variable length packets.
 - b It uses 53 byte units of data called cells.
 - c It can transfer video, voice, and data simultaneously.
 - d It uses predefined circuits.
 - e It is less expensive than other WAN technologies.
- 9 Which of the following technologies is capable of transmission rates of up to 2.488 Gbps?
- a ADSL
 - b T-1
 - c ISDN BRI
 - d OC-SONET
- 10 Which of the following is true of NAT? (Select all that apply.)
- a NAT requires that each computer on the internal network have a registered public IP address.
 - b NAT is compatible with all applications.
 - c NAT is incompatible with technologies that encrypt IP data.
 - d NAT uses a table to map private internal IP addresses to one or more external public addresses.



Physical Components of the Network

Now that you have an understanding of the basic concepts involved in computer networking, we turn our attention to something more concrete: the physical components used to link PCs so that they can share resources.

The components required to link the computers in a network can be as simple and inexpensive as a few low-cost network interface cards (NICs) and a length of Ethernet cable. On the other hand, your network design can be complex enough to require the services of a *network architect* to designate the necessary devices, with a budget running to six figures or beyond.

In this chapter, we discuss some of the most common network hardware devices and networking media, including NICs, cable and wireless media, and connectivity devices.

NICs

The most basic piece of hardware required to network computers is the NIC, also called a network adapter or network card. NICs come in several varieties, which will be discussed in the section "Selecting a NIC."

The NIC is generally referred to as a physical layer device, and the NIC drivers (the software that interfaces between the NIC and the computer operating system) work at the data link layer of the OSI model.

In the following sections, we examine the role of the NIC in network communications, and you will learn how to select the appropriate network adapter and how NICs are configured and used to send and receive data on the network.

The Role of the NIC in Network Communications

Some sort of network interface is always required to communicate over a network. When you connect to a network remotely over analog phone lines, the modem is your network interface and serves the function that a NIC serves on a local network.

The NIC is the basic hardware component of network communications. It translates the *parallel* signal produced by the computer into the *serial* format that is sent over the network

cable. The 1s and 0s of binary communication are turned into electrical impulses, pulses of light, radio waves, or whatever signaling scheme is used by the network media.

An important part of the network interface is the *transceiver*. Some NICs, such as those made for 10Base2 and 10BaseT networks, have the transceiver built onto the card itself. Others, such as those made for 10Base5 networks, have an attachment unit interface (AUI) connector by which a cable is attached to an external transceiver. The transceiver, as its name indicates, sends and receives signals.

Along with preparing the data to go onto the network media, the NIC is responsible for controlling the flow of data between computers and media and for receiving incoming data.

Selecting a NIC

When selecting a NIC for a computer, you should consider the following:

- **Network architecture**—The NIC should be made to work with the existing transmission technology. However, it is possible to use a *media filter* that enables you, for example, to use a Token Ring card on an Ethernet network. Refer to Chapter 6, “WAN Links,” for more information on the popular LAN architectures.
- **Media type**—Ethernet can be run over thick coax (10Base5), thin coax (10Base2), or twisted-pair cable (10BaseT). The connector on the card must match the connector on the cable. Fiber-optic and wireless NICs are also available. We cover cable and wireless media later in this chapter, in the section “Network Media.”
- **Data transfer speed**—If you have 100-Mbps hubs, a 10-Mbps network card does not work. However, you can get dual-speed components (that is, hubs or NICs) that work at either 10 Mbps or 100 Mbps. Likewise, a 4-Mbps Token Ring card does not work on a 16-Mbps network. However, a 16-Mbps card *does* work on the slower network, but its speed drops to 4-Mbps.
- **Available bus type**—Do you have a free ISA or PCI slot? Is the computer a laptop that must use a PCMCIA (that is, a PC card) interface? Do you need a special card that can connect through a serial or SCSI port? Note that if you have both ISA and PCI slots free, the PCI bus is faster.

PC Bus Types

The *data bus* is a transmission path on the computer’s motherboard. Signals are picked up or delivered at each device that is attached to this path or line. There are several different bus architectures for which NICs are made:

- **ISA (Industry Standard Architecture)**—A 16-bit expansion slot on the motherboard. The plastic surrounding the slot is usually black, and the slot (actually two slots, one behind the other) is longer than the PCI slot.

- **EISA (Extended ISA)**—A 32-bit architecture compatible with ISA. Most modern “ISA” slots are actually EISA.
- **PCI (Peripheral Component Interconnect)**—A 32-bit bus used in modern computers (that is, Pentium-class PCs and above and the Macintosh). PCI architecture supports Plug and Play (PnP) devices. The slots on the motherboard are usually surrounded by beige or light-colored plastic and are shorter than ISA/EISA slots.
- **MCA (Micro Channel Architecture)**—A proprietary bus used on IBM computers, which can function as 16- or 32-bit. It is rarely seen today.
- **PCMCIA (Personal Computer Memory Card International Association; also called the PC card)**—An input/output (I/O) bus that uses devices the size and shape of credit cards. Usually used in laptop or notebook computers, PC card adapters for desktop computers are also available.

There are also NICs that connect through the SCSI bus. A few USB (universal serial bus) network cards are available, but these are not common. Also note that some computer motherboards come with an *onboard NIC*, that is, the network interface card is integrated into the motherboard.

- **Operating system**—You must ensure that the manufacturer of the NIC makes drivers for the operating system you are using. If you are using one of the Windows operating systems, check the Microsoft Hardware Compatibility List (HCL) at www.microsoft.com for a list of NICs that have been tested and found to work with the operating system. We discuss the popular desktop operating systems in Chapter 12, “Desktop Operating Systems.”

Configuring and Using a NIC

Configuring a NIC is similar to the process of configuring a modem, as discussed in Chapter 6. Remember that the NIC and modem perform essentially the same basic function.

As with the modem, you might need to set the following:


- IRQ
- I/O address
- Memory address

You can set these through dip switches, jumpers, or software configuration programs that come with the card. PnP NICs can be detected and automatically configured by a PnP operating system, as long as the computer’s BIOS also supports PnP.

Additional configuration parameters may be required if the NIC is a *combo card*, which can be used with different media types. A combo card typically has a BNC connector for thin coax cable, an RJ-45 connector ("RJ" stands for registered jack) for twisted-pair cable, and sometimes, an AUI connector for thick coax cable as well. This means that the same card works on 10Base2, 10Base5, and 10BaseT networks. However, only one media type can be used at a time. You might have to manually select the media type, or the card might autodetect the type of media that is attached.

Figure 7-1 shows a three-way combo card.

Figure 7-1 A combo NIC can be used with 10Base2, 10Base5, and 10BaseT networks.



Configuring the IRQ

Typically, a NIC uses IRQ 3 or IRQ 5. Many cards are set to use one of these by default. If you have another device using the predefined IRQ, you need to change the settings for one of the devices so that they do not conflict.

It is a good idea to memorize the following table, which shows the standard default IRQ settings for a typical PC. Note that in early PCs, only the first eight IRQs were available.

IRQ 2 is used as an IRQ controller, to “redirect” an interrupt request to IRQ 9, through which the high-numbered IRQs can be accessed.

Table 7-1 *Standard Default IRQ Settings for PCs*

Device	Default IRQ Used
VGA graphics adapter	2 (9)
COM 2 and COM 4 (secondary serial port)	3*
COM 1 and COM 3 (primary serial port)	4
Secondary parallel port (LPT2) or sound card	5*
Floppy disk controller	6
Primary parallel port (LPT1)	7
Real-time clock	8
Primary SCSI controller	10*
Secondary SCSI controller	11*
PS/2 mouse	12
Math coprocessor	13
Primary IDE controller	14
Secondary IDE controller	15*

* These devices are optional and might not be present on many computers. If these devices are not present, the corresponding IRQs are free. For example, if the computer has no SCSI controllers, IRQs 10 and 11 are probably free and can be assigned to the NIC.

Configuring the I/O Port and Memory Address

In addition to a unique IRQ, each device must have a different I/O port setting. The *I/O port* is a channel through which data is transferred between the hardware device and the processor. Port numbers are designated by hexadecimal numbers. The following ports are usually available for use by the NIC:

- 300 to 30F
- 310 to 31F

Others may also be available. Generally, the default I/O port settings (to which the NIC was set at the factory) work.

A *memory address* is a location in RAM that is used for storage of incoming and outgoing data. Some NICs do not use the computer’s memory for this purpose and thus do not have a memory address setting. Otherwise, the default factory settings generally work.

Network Media

The network *media* is the means by which the signals travel from one networked device to another. The most common media type is cable, but newer wireless media are now available as well (for example, radio waves, laser and infrared beams, satellite, and microwaves).

In this section, we first discuss different cable types commonly used to join networked computers, and then we look at wireless options.

Cable Types

Most networks today are cabled, and the most common types of cable in use are coaxial, twisted-pair, and fiber-optic.

Coaxial Cable

Coaxial cable is familiar to most persons who have cable TV. It has a copper core (which can be either stranded wire or solid copper). The signal travels on the copper, which is wrapped in insulation. Surrounding the insulation is another conductor of metal foil or braid. This outer conductor runs the length of the cable, hence the name coaxial, because two physical channels, one carrying the signal and the other serving as a ground, run together (“co”) along the same *axis* (“axial”). The outer conductor acts as a shield against outside electromagnetic interference (EMI). These components are then wrapped in an outer shielding of plastic, rubber, or—in the case of *plenum grade* cable—Teflon or other fire-resistant material.

NOTE

Plenum-grade cable is required by most local building and fire codes when cable is installed in the *plenum*, the space between a false ceiling and the floor above it, because standard grade coaxial cable is covered with poly-vinyl chloride (PVC), which emits toxic gases when burned.

Because of its thick insulation and shielding, coax is less vulnerable to outside EMI than is twisted pair.

There are thousands of different types and grades of coax cable, as a glance at any cable manufacturer’s catalog shows. Many are used for special-purpose networks, such as connecting scientific instruments or other dedicated devices. The types discussed in the

following sections and summarized in Table 7-2 are the only ones with which you must be familiar for most PC LAN networking tasks.

Table 7-2 *Coax Cable Types and Characteristics*

Designation	Common Name	Description	Network Use
RG-8, RG-11	Thicknet	One-half inch diameter thick coaxial cable	10Base5
RG-58 A/U	Thinnet	One-quarter inch diameter thin coaxial cable	10Base2
RG-58 C/U	Thinnet(military spec)	One-quarter inch diameter thin coaxial cable	10Base2 (military use)
RG-62	ARCnet	Thin coaxial cable	ARCnet networks

In early implementations of Ethernet, coaxial cable was the most popular type. However, twisted-pair coaxial cable has surpassed it in popularity. Coaxial cable for Ethernet networking comes in two basic types: thin coax and thick coax, also called *thinnet* and *thicknet*.

NOTE Cable TV (CATV) coax looks very much like thinnet, but they are not interchangeable. CATV cable is RG-59, 75-ohm cable.

Thin Coax

Thin coax cable is approximately one-quarter inch in diameter and more flexible than thicknet. It is used in Ethernet 10Base2 networks, and it can transmit signals for approximately 185 meters without suffering from attenuation (that is, signal weakening).

Thin coax cable is sometimes called RG-58 cabling. Its *impedance* (that is, its resistance to current flowing through the wire) is 50 ohms.

Basic Electronics Terminology

Electrical current is a flow of electrons, or electron-deficient atoms, that carry an electrical charge. An *ampere* (amp) is the basic unit used in the measurement of electrical current.

Resistance (also called impedance) is a way of measuring the amount of opposition to the flow of electrical current. Different substances offer different levels of resistance.

Resistance is measured in *ohms*, named after German scientist Georg Simon Ohm, who

continues

postulated the mathematical relationship among electrical current, resistance, and voltage (Ohm's Law). The symbol for ohm is the Greek letter omega (Ω).

Voltage or *electromotive force* (EMF) is also called the electric potential. It is a measurement of the force that causes movement of electrons in a conductor (a medium that conducts, or transmits, an electric charge).


Frequency is the measure of how often a periodic event takes place (for example, a signal completing a cycle). *Hertz* is a unit used to measure frequency.

Coax cable is standardized throughout the industry according to RG (Registered Grade) specifications. Only a few of the hundreds of coax grades are used in PC networking.

10Base2 networks use type RG-58 A/U thin coax. This type has a stranded wire core. The military specification for the same cable type is RG-58 C/U. The military grade cable has a slightly thicker outer covering, making it more resistant to heat, cold, liquid, and other elements.

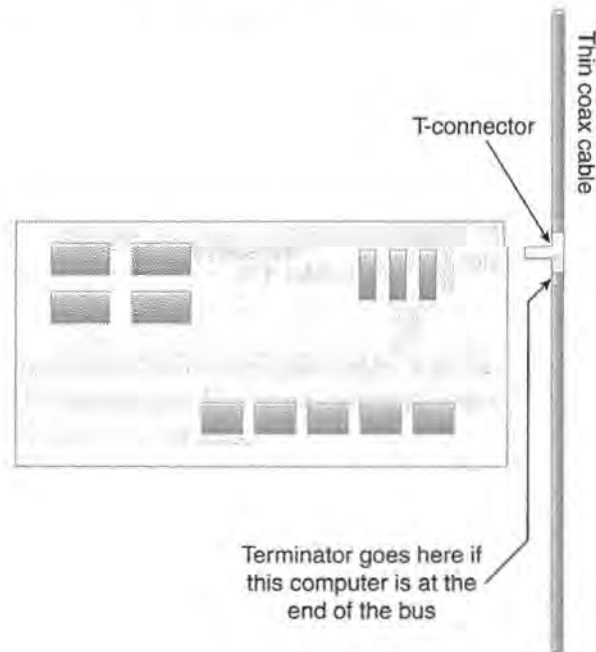
RG-58 /U is similar in appearance to RG-58 A/U, but the copper core is solid. It does not meet 10Base2 specifications, and you should *not* use RG-58 /U on a computer network. It is used only for antenna wiring and other purposes. See Figure 7-2 for an illustration of the difference between the two.

Figure 7-2 *The difference between stranded and solid core copper cable.*



Thinnet cable connects to the network card through a BNC T-connector that attaches to the NIC's BNC connector, as shown in Figure 7-3.

Figure 7-3 Thin coax connects to the NIC with a BNC T-connector.




A terminator is required at each end of a thinnet network. The terminator connects to the unused side of the T-connector.

Thick Coax

Thick coaxial cable, also called thicknet, is approximately twice the diameter of thinnet—about one-half inch. Thicknet was the original cable type used in early Ethernet networks, and thus is referred to as *standard Ethernet*.

Thick coax is type RG-8 or RG-11 cable. Because of the thicker core, it can transmit signals for a longer distance—500 meters—without attenuation. It is, however, more expensive and more difficult to work with than thin coax because it is less flexible and it uses an external transceiver that must be connected to the cable with a device called a *vampire tap*. The tap drills into the core of the cable. Making the connection is more complex than using twist-on BNC connectors on thinnet, as shown in Figure 7-4.

Figure 7-4 A 10Base5 network uses an external transceiver, connected to thick coax by vampire taps.



NOTE

Another type of coax cable, commonly used in the past in ARCnet networks, is RG-62 90-ohm cabling. The ARCnet architecture has become less popular in recent years, as discussed in Chapter 5, "LAN Links."


Twisted-Pair Cable

Many people today, when they use the term *Ethernet cable*, are referring to the twisted-pair cabling used in a huge number of Ethernet networks. It is called *twisted pair* because, inside the cable's outer sheath, pairs of insulated copper wires are twisted around one another to prevent *crosstalk* (that is, the "leaking" of signal from one wire to another). As the number of twists per foot increases, so does the degree of protection against crosstalk.

Telephone companies use twisted-pair cable for internal wiring; this contributes to its popularity because most buildings are already wired with twisted-pair cable for telephone systems.

Twisted-pair cable comes in two basic types: unshielded twisted-pair (UTP) cable and shielded twisted-pair (STP) cable, as shown in Figure 7-5. Telephone communications and most Ethernet networks use UTP. STP is used in Token Ring and AppleTalk networks.

Figure 7-5 *The difference between STP and UTP.*



When shielding is added to twisted-pair cable, it reduces the effect of outside electromagnetic interference. However, it increases attenuation and can affect the resistance of the wire and cause loss of data.

UTP Cable

UTP is a very popular cable type for LANs, for several reasons:

- It is relatively inexpensive.
- It is flexible and easy to work with.
- It uses familiar RJ-45 connectors that look and work like the smaller RJ-11 modular telephone connectors.
- It is used in a star topology, which offers advantages that are discussed in Chapter 2, “Categorizing Networks.”

UTP cabling is rated by category, according to its use and data transmission speed. Refer to Table 2-3 in Chapter 2 for a listing of UTP categories and the uses and transmission speeds of each.

NOTE

UTP is the telephone wiring that is inside walls and that connects to an outlet called an RJ-11 receptacle (jack). The common telephone wire that runs from the wall to the phone is *untwisted* flat copper pair, often called “silver satin” because of its shiny gray color.

Like coax, UTP comes in both standard and plenum-grade versions.

STP Cable

Shielding is made of woven copper or foil, and it wraps the insulated copper pairs inside the outer covering. This reduces the effect of EMI, but causes STP to be more expensive than UTP and introduces other problems, as discussed in previous sections of this chapter.

STP is often used for AppleTalk and IBM Token Ring networks.

Fiber-Optic Cable

Fiber-optic cable (also called *optical fiber*, *optical cable*, or *fiber*) is a newer, faster, but relatively expensive transmission medium that is growing in popularity as high-bandwidth applications become more common. Although often implemented at 100 Mbps, optical fiber is capable of speeds of 1 Gbps or more.

NOTE

Lucent Technologies has documented data transmission speeds of over 3 *terabits* per second in the laboratory, using multiple lasers with fiber-optic cable. A terabit is one trillion bits.

In lieu of copper, fiber-optic cable uses tiny strands of glass or plastic through which the signal is transmitted in the form of light pulses.

Fiber-optic cable has several advantages, in addition to speed:

- It is more secure than coax and twisted-pair cable because there is no electrical signal that can be tapped.
- It is less susceptible than other cable types to attenuation and thus can span long distances—2000 meters or more.
- It is not vulnerable to outside electrical interference such as EMI and radio frequency interference (RFI).

Fiber-optic cable is costly compared to more traditional cable types. Although the cable itself is somewhat more expensive, the largest expense is labor; fiber-optic cable is much more difficult to work with and requires specially trained technicians to splice the tiny strands of glass or plastic.

Fiber-Optic Mode Types

Fiber-optic cable operates in one of two modes:

- **Single mode**—Single mode is also called *axial* because the light travels down the axis of the cable.

- **Multimode**—In multimode fiber, light waves enter the glass pipe at different angles and travel *nonaxially*, which means that they bounce back and forth off the walls of the glass tube.

Single-mode fiber is faster than multimode (up to 10 Gbps) because of the *dispersion* (scattering or separation of light waves) in multimode caused by the light pulses arriving at the end of the cable at different times. Single mode is typically used for WANs (for example, telephone company switch-to-switch connections). Multimode is often used in LANs.

Fiber-Optic Cable Light Sources

Fiber-optic cable can be categorized by the type of light source used:

- **LED (light emitting diode)**—LED is commonly used with single-mode fiber. It is relatively weak.
- **ILD (injection laser diode)**—ILD emits a strong, intense, narrowly focused light beam. It is commonly used with multimode fiber, which helps counteract multimode's lower performance.

LEDs are widely used in displays for digital clocks, remote controls, and electronic instruments. The LED is a semiconductor device. When a current of electricity passes through it, the LED gives off light. This light is typically red but can range through the spectrum to a blue-violet color depending on the wavelength. LEDs that emit infrared energy are called *infrared emitting diodes (IREDS)*.


ILDs are sometimes used in hand-held laser scanners.

Wireless Media

Wireless technologies are becoming a popular networking alternative. Although wireless transmission methods are often much slower than cabled connections, there are significant advantages to wireless networking under certain circumstances, as discussed in Chapter 6. The same basic concepts that affect wireless WANs also affect wireless LANs.

Typically, a “wireless” network is not completely without cables, but incorporates wireless devices that communicate with a traditional cabled network. Transceivers, which are called *access points*, are used to transmit and receive data between the wireless device or devices and the wired network, as shown in Figure 7-6.

Figure 7-6 *Wireless networking uses transceivers to transmit and receive data without cables.*



NOTE Cisco Systems, in partnership with other companies such as Motorola, EDS, Samsung, and Texas Instruments, is leading the way in developing high-speed, reliable, and widely available wireless technologies.

Wireless LAN communication is a category that includes a variety of transmission methods, including the following:

- Laser
- Infrared
- Radio

We look at some of the characteristics of each in the following sections.

Laser

Laser is an acronym for “light amplification by stimulated emission of radiation.” A laser outputs a coherent electromagnetic energy field, in which all waves are at the same frequency and are aligned in phase. A *phase* is a fraction of a complete cycle that has elapsed and that is measured from a specific point of reference. Different types of lasers produce beams of different wavelengths.

NOTE

Most people think of laser light as a “red dot,” but argon lasers (so called because they use argon gas as the laser medium) produce a blue or green light. Krypton lasers (using gas by the same name) produce red light, and “mixed gas” lasers combine the two gases, which produces red, green, and blue output simultaneously. This simultaneous output results in a white beam.

Lasers are used for many purposes, ranging from cheap pointing devices to highly accurate pistol sights, printers, electronic games, remote control devices, surgical procedures, and network communications.

Laser networking works by using pulses of laser light to represent the data signals. Laser is a *line-of-sight* technology, which means that there must be an unobstructed pathway between the transmitting and receiving devices. The need for this unobstructed pathway is a drawback of laser-based wireless communications.

Infrared

Infrared (IR) technology is familiar to many people because of its use in TV remote control units. It can also be used to network wireless LANs by using cones or beams of light in the IR frequency spectrum to carry the data signal. IR uses very high frequencies that are just below the visible light spectrum.

The Infrared Data Association (IrDA) is an organization that sets standards for IR hardware and software. IrDA-compliant devices are designed so that when a user breaks the infrared connection, the connection is reinstated when the devices move back into IR range.

IR networking requires a transceiver in both communicating devices and might require synchronization software as well. Some operating systems, such as Windows 98 and 2000, have built-in IR support. Transfer speeds range from 4 Mbps to 16 Mbps.

IR, like laser, is normally a line-of-sight technology. However, in some implementations of IR, such as scatter and reflective, the signal can be bounced or redirected. Even in these implementations, however, IR cannot go through opaque objects such as walls.

Disadvantages of IR for networking include the following:

- **Distance limitations**—Although it is possible to implement IR at distances greater than one mile, the more common distance is less than 100 feet.
- **Vulnerability**—Ambient light can cause interference.

Broadband optical telepoint is an IR technology designed to support high-bandwidth multimedia applications.

Radio

We know that telephone lines, originally used to transmit voice, can transmit data as well. Likewise, radio waves—a medium we associate with audio transmission—can be used to carry data signals.

Data transmission over radio can be implemented with a range of technologies. We look at two broad categories in the following sections:

- Narrowband radio
- Spread spectrum radio

Narrowband Radio

Narrowband radio is familiar to most of us. A transmitter sends a signal on a specified frequency, and a receiver tuned to that frequency picks up the signal. This is how broadcast radios, two-way radios, and traditional emergency channel communications work.

When data is transferred over narrowband radio, it is easy for an unauthorized listener to intercept the signals. A more secure and more reliable radio technology, originally developed by the military, is called *spread spectrum*.

Spread-Spectrum Radio

Spread-spectrum radio is a *wideband* technology. Although it is less efficient than narrowband (that is, it uses more bandwidth), it is more secure because it uses multiple frequencies. Narrowband receivers are not able to pick up spread-spectrum transmissions. There are two main types of spread spectrum radio:

- **Frequency hopping spread spectrum (FHSS)**—With FHSS, the transmitter hops from one frequency to another. The receiver must know the frequencies, the pattern, and the timing of the hops. This makes it difficult for an unauthorized person to intercept the signal.
- **Direct sequence spread spectrum (DSSS)**—This technique uses special encoding (calling *chipping*) that creates a redundant bit pattern for each bit of transmitted data. This provides fault tolerance because if some bits are damaged during transmission, the original data can still be recovered without retransmission.

Network Connectivity Devices

Connectivity device is a general term that includes the simple and complex devices used to connect one part of a network to another. In most cases, this means two or more lengths of cable are connected to a device.

In the next sections of this chapter, we examine three types of connectivity devices:

- Simple connectors
- Complex connectors
- Segmenting and subnetting devices

Simple Connectors

Simple connectivity devices are those that provide only a connection point and that do not amplify or otherwise modify the signal. These include the following:

- BNC connectors
- RJ connectors
- Fiber-optic connectors
- Patch panels
- Passive hubs

The following sections describe these connectors in more detail.

BNC Connectors

10Base2 (that is, thin coax) networks use BNC connectors to connect the NIC to the cable. A BNC connector is a small cylindrical device with a pin that connects to the conduction wire in the cable. The connector locks into place by the twisting of an outer ring.

BNC devices include the following:

- **BNC T-connector**—With this connector, the stem of the T attaches to the NIC, and a piece of cable attaches to each side of the top bar. If only one cable is to be connected, a terminator must be connected to the other side of the T-connector.
- **BNC barrel connector**—This connector is a straight cylindrical unit to which a cable attaches at each end, thus enabling you to join two pieces of cable to increase the total cable length.

WARNING The use of barrel connectors should be kept to a minimum because of signal loss that can occur at each connection point.

The BNC *terminator* is a 50-ohm termination device, which is installed at each end of a coax bus. The terminator prevents a signal from bouncing back when it reaches the end of the cable, which causes interference. Both ends of the cable should be terminated, and one

end should be *grounded* by attaching a conductor such as a wire to a position of zero electrical potential. The term “ground” is used because, in many cases, the conductor is physically connected to the actual ground (the earth).

RJ Connectors

RJ (registered jack) connectors are so called because they are registered with the Federal Communications Commission (FCC). RJ connectors consist of a *plug* and a *receptacle*. The receptacle is sometimes referred to as the *jack*.

Ordinary analog telephone wires in the United States generally use RJ-11 connectors. These are modular plugs connected to “silver satin,” the familiar flat gray phone cord that runs from the wall outlet to the telephone. Inside the wall, the jack connects to UTP (usually Cat 3 or 5 in modern buildings). Modems have RJ-11 jacks.

RJ-45 connectors are used on UTP cable designed for Ethernet networks. The RJ-45 plug and receptacle look like the RJ-11, but are slightly larger. A special crimping tool is used to attach the wire pairs to the RJ plug.

There are numerous RJ designations, but RJ-11 and RJ-45 are the ones commonly encountered in computer networking.

Fiber-Optic Connectors

Connectors for optical cable are the most difficult to install because each individual strand of glass or plastic must be precisely aligned. Once they are aligned, the cable is attached to the connector with hot melt glue, epoxy, or anaerobic adhesive.

Common fiber-optic connectors include the following:

- **SC**—A push-pull type
- **ST**—A keyed, bayonet type
- **FC**—A keyed, threaded-lock type
- **SMA**—A threaded type
- **FSD**—A fixed shroud device, used with FDDI

Patch Panels and Passive Hubs

A *patch panel* is a connection and distribution point used to organize cables that come together at a central location. This works somewhat like early telephone switchboards, in which a connection was completed by plugging a wire into a specific jack.

For a star topology network, cables from computers at various locations in the building come into a wiring closet, which is where the patch panel is located. The cables are

connected to a punch down block on the back of the panel. On the front are RJ-45 jacks, from which patch cables run from the panel to the hub.

A patch panel serves the same function as a passive hub. A passive hub is a central connection device where the cables from computers and other network devices meet. It contains no electronic parts and does not require electrical power. It does not regenerate the signal, but merely sends it out through all ports of the hub.

Complex Connectors

Simple connectivity devices merely connect cables; complex devices do more. For example, they might strengthen the signal before passing it on and even convert the signal from one media type to another.

In the following sections, we discuss the following types of complex connectors:

- Media converters
- Repeaters
- Active/intelligent hubs

Media Converters

Media converters are also called *media adapters* or *media translators*. They are used to convert one segment type to another; for example, they can convert 10Base2 to 10BaseT, 100BaseT Ethernet to fiber optics, or Token Ring to fiber optics.

Repeaters

A repeater connects two network segments or lengths of cable. Unlike a barrel connector, however, it doesn't just pass the signal on from one cable to the next—it *regenerates* the signal. Thus a signal that has weakened because of attenuation is strengthened, and the effective distance of the cable is increased.

NOTE

Repeaters can be used to connect different media types, such as a 10Base2 segment to a 10Base5 segment. Both architectures are Ethernet and use the same media access method. Repeaters cannot be used to connect segments using different architectures or access methods. For instance, you cannot connect an Ethernet segment (which uses CSMA/CD) and a Token Ring segment (which uses token passing).

Repeaters do not filter the data that passes through them. They regenerate *all* signals, including broadcast messages, noise, and interference, and pass them on. Repeaters operate at the physical layer of the OSI reference model.

Active and Intelligent Hubs

Active hubs are also called *multiport repeaters* because they have multiple ports (like a passive hub), and they regenerate the signal coming into one port before sending it back out the other port (like a repeater). Active hubs require electrical power.

The intelligent hub is a special type of active hub. It not only regenerates the signal but also has an onboard processor that enables you to perform diagnostics and detect if there is a problem with a particular port. Hubs operate at the physical layer of the OSI reference model.

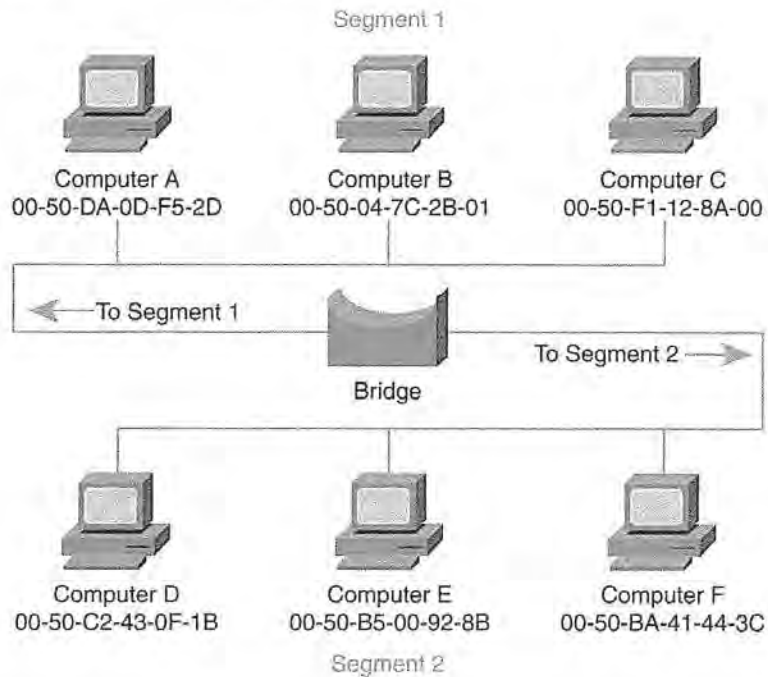
Segmenting and Subnetting Devices

Segmenting and subnetting devices are the most complex of the network connectivity devices. Although the terms are sometimes used interchangeably, *segmenting* the network refers to dividing it into segments that are still part of the same network. *Subnetting* goes a step further and divides the network into separate networks called *subnetworks* (based on network address information). We discuss subnetting in more detail in Chapter 8, “Networking Protocols and Services,” and Chapter 9, “The Widest Area Network: The Global Internet.”

Bridges

A traditional bridge (also called a *simple transparent bridge*) joins two network segments and performs *filtering* of traffic, based on the media access control (MAC) address on the packet. When used properly, this enables you to reduce congestion. The network is divided into two segments, with the bridge in between, as shown in Figure 7-7. MAC addresses are indicated as hexadecimal numbers.

Figure 7-7 A bridge divides the network into two segments.



The bridge builds an address table by executing the following steps:

- 1 When a packet is sent on the network, the bridge checks the source and destination address (that is, the MAC address). The table tells the bridge on which segment (that is, which side of the bridge) each address is located.
- 2 If the destination address of a packet is not in the bridge's table, the bridge forwards it to both segments. If the source address is not in the table, the bridge adds it to the table.
- 3 If the destination address is in the table, the bridge forwards the packet to the appropriate segment *unless* the source and destination computer are on the same segment.
- 4 If the table shows that the source and destination are on the same segment, the bridge does not forward the packet.

NOTE

The address table built by the bridge is called a *routing table* because it is used to determine to which side the packets should be routed. Don't confuse this with the routing table used by a router. The bridge's routing table uses hardware addresses, but the router's table is based on higher-level IP addresses.

In the example shown in Figure 7-7, if Computer B sends a message to Computer F, the bridge forwards the packet to Segment 2. It does this regardless of whether Computer F's address is in the table because the packet is forwarded to all segments if there is no address entry.

However, if Computer B sends a message to Computer A, the bridge checks its table. If Computer A's address is not there, the packet goes across the bridge to the other segment. If Computer A's address has been entered (because Computer A previously sent a message), the message does *not* go across the bridge. You can see how this reduces unnecessary traffic on Segment 2.

The bridge is called *transparent* because the computers on an Ethernet network are not aware of its presence.

Bridges forward broadcast messages, which are those that are addressed to the hardware *broadcast address* (FF-FF-FF-FF-FF-FF).

Translation and Encapsulation Bridges

Unlike repeaters, some bridges can connect network segments using different media access methods (for example, Ethernet and FDDI), as long as they use the same network protocol (for example, TCP/IP). These are called *translation bridges* or *encapsulation bridges*.

The translation bridge translates the Ethernet addresses into FDDI addresses. You can also bridge unlike networks by using encapsulation bridging, in which the Ethernet frame is *encapsulated*, or wrapped, inside a FDDI frame.

A *source routing bridge* is a special type of bridge used on Token Ring networks. It is unlike standard transparent bridges because it depends on the host computer to make the routing decision.

Bridges operate at the data link layer of the OSI reference model. Thus, *nonroutable* protocols, such as NetBEUI, can cross bridges. (You learn more about routable and nonroutable protocols in Chapter 8.) Like repeaters, bridges regenerate data, but they do so at the *packet* level.

A network can have more than one bridge. This provides fault tolerance, but can lead to a *bridging loop* problem, which occurs when there are multiple paths between two points, and packets end up going around in circles. This creates unnecessary traffic (see Figure 7-8).

The *Spanning-Tree Algorithm (STA)* was developed to solve the bridging loop problem. STA creates a subset of bridged links that eliminates looping. The technical details of bridge looping and STA are beyond the scope of this book, but excellent resources are listed at the end of this chapter.


Routers

As we work our way up the OSI reference model, we encounter the connectivity device that works at the network layer: the *router*. Routers connect separate networks to one another. This can occur within a LAN (in which case the individual networks are called *subnets*) or between unrelated networks in a WAN such as the global Internet.

Like a bridge, a router filters traffic. Unlike a bridge, it does so using the logical network address (IP or IPX address) instead of the physical hardware address. Routers are more intelligent than bridges; they make complex decisions by selecting the best route to a given destination from among multiple paths.

NOTE

Dedicated routing devices are actually special-purpose computers; they contain microprocessors and run their own operating system. PCs can also be configured to act as routers if the operating system supports IP or IPX forwarding.

Figure 7-8 Multiple paths between bridges can result in a bridging loop.

Routers maintain routing tables that contain the network addresses of other routers. A router must have at least two network interfaces because it serves as a *gateway* from one network to another. The address of the router's interface that serves a particular subnet is called that subnet's *default gateway*. (Note that the term "gateway" is also used to describe software and devices that function at the higher OSI layers to translate between protocols.)

Router Functionality

Routers can be used to join multiple networks into one larger one or to separate a large network into several smaller ones. In Chapter 9, we discuss why networks should be subnetted.

When a data packet travels from one router to another, the data link layer headers (that is, the pieces of addressing information) are stripped off and recreated. This enables routers to exchange packets between unlike networks, such as Ethernet and Token Ring. This process does require overhead, however, which makes the network performance slower than it would be with lower-level devices.

When multiple paths exist on a network, a bridge chooses one and always uses that path to reach a specific destination. Routers consider all available paths for each packet sent and make the decision on a packet-by-packet basis. Thus, if one route is very busy at that time, the router chooses another, more efficient one.

Another advantage of the router is its capability to filter both Layer 2 and Layer 3 broadcast packets. By default, routers do *not* forward messages that were sent to the broadcast IP address (255.255.255.255) across the router. This reduces network traffic significantly and prevents the propagation of *broadcast storms*. A broadcast storm occurs when there are so many broadcasts that the network cannot function properly.

Routable Protocols Versus Routing Protocols

Routers work only with *routable* protocols, including IP, IPX, OSI, XNS, DECnet, and DDP. A nonroutable protocol, such as NetBEUI, does not use an addressing scheme that enables the router to identify the network; thus, the protocol cannot be routed.

It is important to distinguish between routable protocols (discussed previously) and *routing* protocols. The latter are used by the router and are required for *dynamic routing*.

In the following sections of this chapter, we compare routing types and routing protocols, including the following pairings:

- Static versus dynamic routing
- Interior versus exterior routing protocols

Static and Dynamic Routing There are two basic ways of routing:

- **Static routing**—This requires that an administrator manually enter addresses into the routing table and keep that table updated.
- **Dynamic routing**—This uses protocols such as Routing Information Protocol (RIP), Open Shortest Path First (OSPF), or NetWare Link Services Protocol (NLSP) to enable routers to automatically and dynamically exchange routing table information with one another.

You will learn more about dynamic routing, and these protocols, in Chapter 9.

Interior and Exterior Routing Protocols Routing protocols are classified as either *interior gateway protocols (IGP)* or *exterior gateway protocols (EGP)*. These are discussed in more detail in Chapter 9.

Routing protocols that operate within an *autonomous system*—that is, a network under the control of a particular company or organization—are IGPs.

Brouters

A brouter can function as either a bridge or a router, depending on the network transport protocol in use. A brouter acts like a bridge for messages sent with NetBEUI or other nonroutable protocols, but it provides the functionality of a router for routable protocols such as TCP/IP. Modern routers are capable of bridging and routing.

Switches

The basic functionality of a switch is deceptively simple: choosing a path across which to send data to its destination.

Ethernet switches are becoming a popular connectivity solution, and for good reason. They increase performance (speed) and are relatively inexpensive.

Switches use one of two switching schemes:

- **Cut-through switching**—The switch starts forwarding the packet to its destination before it has received the entire packet. This method is faster, but it can result in bad packets getting through.
- **Store-and-forward switching**—The switch does not send the packet until it has been completely received and its integrity has been checked. This is slower, but more reliable.

There are different types of switches. Switches are sometimes categorized based on the layer of the OSI reference model at which they function. In the next sections, we discuss the differences between Layer 2, Layer 3, and Layer 4 switching.

Layer 2 Switching

Standard Layer 2 switches act like hubs—with an important difference. Where a hub sends messages out all its ports, a switch (referred to as a *switching hub*) is “smart” enough to determine which port is connected to the computer to which the message is addressed and to send it only to that port. This has several positive effects:

- The overall amount of unnecessary network traffic is reduced, which decreases congestion.
- Separate collision domains are created, which prevents data collisions that slow performance and that require resending of messages.
- Security is increased because messages are not going out all ports. Messages that are going out on all ports are easier to intercept.

These switching hubs are also called *port switches* because a computer or network device connects to each port. Each device has its own dedicated pathway to the switch.

Another type of switch is the *segment switch*, which enables you to connect an entire network segment to each port.

Switches can be used to create *virtual LANs (VLANs)*, which divide the physical network connected to a switch into multiple logical networks. This can increase both performance and security.

Layer 3 Switching

Layer 3 switches, as the name implies, operate at the network layer of the OSI reference model. This device was first developed by 3Com in 1992, when it started to integrate its switching and routing devices to reduce the number of devices that needed to be managed.

The most important thing to understand about Layer 3 switches is that they are routers, but of a special type. A Layer 3 switch (or *switched router*, as it is sometimes called) performs the same functions as a dedicated router and uses routing protocols, such as RIP and OSPF.

The difference is that Layer 3 switches perform the functions that a Layer 2 switch performs. The switch uses a hardware-based architecture to apply policies based on network layer information in the packet header.

Layer 3 switches are generally easier to set up and configure than are routers and can be used in most situations (within a local network) where a router can be used. Surprisingly, despite the added functionality and ease of use, Layer 3 switches are generally less expensive than comparable routers.

Layer 4 Switching

Recent enhancements to Layer 3 switches enable them to use information, such as port numbers, from the TCP and UDP headers. These enhancements are referred to by some as *Layer 4 switching* because TCP and UDP operate at the transport layer (Layer 4) of the OSI reference model. Despite the name, these switches are often capable of using information at higher layers as well.

An important use of Layer 4 switching is providing access control list (ACL) filtering for security purposes. Although traditional routers can “see” Layer 4 information—and some routers (such as the Cisco 7500) can perform Layer 4 functions—enabling ACLs results in a significant performance hit. Because the packet processing is done in the hardware with Layer 4 switches, there is not a corresponding reduction in performance.

Layer 4 switches have the capability to manage allocation of bandwidth for quality of service (QoS) implementations and to perform load balancing.

Layer 4 switching functions are available on many modern routers.

Summary

In this chapter, we discussed a broad range of network hardware devices. You learned how to select, install, configure, and troubleshoot various types of NICs.

We also discussed the many choices available in selecting networking media, both cabled and wireless. Specifically, we addressed the characteristics of the three common cable types: coax cable (thin and thick), twisted-pair cable (shielded and unshielded), and fiber-optic cable. You learned about some of the advantages and disadvantages of each, and how they can be used in different LAN configurations.

Then we examined the popular devices that can be used to connect segments or subnetworks. We first looked at simple, nonelectronic devices, such as cable connectors, patch panels, and passive hubs. Then we took a closer look at more complex devices such as repeaters, active hubs, and intelligent hubs. Finally, we learned about segmenting and subnetting solutions such as bridges, routers, and switches.

Now that you have a good understanding of the hardware components that can be involved in the typical networking scenario, in the next chapter we take an in-depth look at networking protocols and services on which the network runs.

Further Reading

An excellent tutorial on wireless networking is available on the Web at www.proxim.com/wireless/whiteppr/whatwlan.shtml.

For more detailed information on bridges, bridge loops, and the Spanning-Tree Algorithm, see www.officewizard.com/books/network/ch07.htm.

An excellent explanation of Layer 3 switching can be found at www.3com.com/technology/tech_net/white_papers/500660.html.

Review Questions

The following questions test your knowledge of the material covered in this chapter. Be sure to read each question carefully and select the *best* correct answer or answers.

- 1 What is the most basic piece of hardware required to network computers?
 - a Router
 - b Cable
 - c NIC
 - d Hub

- 2 Which type of network uses an external transceiver?
 - a 10Base2
 - b 10Base5
 - c 10BaseT
 - d All of the above
- 3 If you want to install a NIC in a computer that has a PS/2 mouse, keyboard, a printer, and two IDE controllers, all using default IRQs, and a sound card set to IRQ 10, which of the following IRQs should be available for the NIC?
 - a 5
 - b 7
 - c 12
 - d 14
- 4 What type of cable is required by most fire codes when you install cable in the space between the false ceiling and the floor above?
 - a Cat 5
 - b Thick coax
 - c Military grade
 - d Plenum grade
- 5 Which cable type is specified for use on the typical 10Base2 network?
 - a RG-58 /U
 - b RG-58 A/U
 - c RG-58 C/U
 - d RG-59
- 6 What cable type is associated with the use of a vampire tap?
 - a UTP
 - b STP
 - c Thick coax
 - d Thin coax

- 7 What is the leaking of signal from one wire to another, which can be diminished by twisting wires around one another?
 - a RFI
 - b Attenuation
 - c Distortion
 - d Crosstalk
- 8 What connector type is generally used with UTP Cat 5 cabling on Ethernet networks?
 - a RJ-45
 - b RJ-11
 - c BNC connector
 - d F connector
- 9 Which of the following is true of fiber-optic cable? (Select all that apply.)
 - a Fiber-optic cable is the easiest of all cable types to work with and to install.
 - b Fiber-optic cable provides better security than does copper cable.
 - c Fiber-optic cable is more susceptible to attenuation than is copper cable.
 - d Single-mode fiber-optic cable is faster than multimode fiber-optic cable.
- 10 Which of the following is used to connect two segments of network cable and to regenerate the signal as it passes through, thus countering the effects of attenuation?
 - a Barrel connector
 - b Patch panel
 - c Repeater
 - d Passive hub



Networking Protocols and Services

If the network's hardware is analogous to its bones and organs, and the signals that run through it can be compared to its lifeblood, we can think of the protocols and services on which the network runs as the mind of the network. The protocols tell the network how to perform its functions. They control it just as the mind controls the body.

The network protocols are sets of rules—the logic by which the network operates. Network *services*, such as name resolution or address allocation services, perform specific functions and control particular tasks. This is somewhat similar to the way different areas of the brain govern hearing, seeing, breathing, speech, and so on.

There are many different types of computer protocols. Each operates at a different layer of the OSI reference model. In the context of PC networking, the term “protocol” is often used to identify the *network/transport* protocols, which are those that work at Layers 3 and 4 of the OSI reference model. Networked computers must use a common protocol to communicate.

We discuss three network/transport protocols in this chapter:

- The NetBIOS Extended User Interface (NetBEUI)
- The Internet Packet Exchange/Sequenced Packet Exchange (IPX/SPX)
- The Transmission Control Protocol/Internet Protocol (TCP/IP)

These are the three standard stacks supported by many popular PC operating systems. Each has advantages and disadvantages, depending on the LAN environment.

In this chapter, we also examine networking services such as Domain Name System (DNS), Windows Internet Name Service (WINS), and Dynamic Host Configuration Protocol (DHCP) that work with TCP/IP to enhance its functionality.

NetBIOS/NetBEUI

You might hear references to NetBIOS and NetBEUI that imply they are the same thing. In fact, at one time this was true. NetBIOS (Network Basic Input/Output System) was developed by IBM and adopted by Microsoft for early LAN communications. At that time,

the term referred to both the application programming interface (API) and the network/transport protocol stack.

NOTE

A *protocol stack* is a group of two or more protocols that work together, with each operating at a different layer of the OSI reference model.

Subsequently, these components were split into NetBIOS (referring to the API) and NetBEUI (which encompasses the network/transport layer protocols). NetBIOS does not provide a frame or data format for network transmission; NetBEUI does.

NetBIOS: The API

NetBIOS can run over NetBEUI, IPX/SPX or TCP/IP. NetBIOS enables applications to deal with a common programming interface so that information can be shared over different lower-level protocols.

Operating at the upper layers (that is, the application layer of the Department of Defense (DoD) model and the session layer of the OSI reference model), NetBIOS provides for two communication modes: session mode and datagram mode.

When running in session mode, NetBIOS enables the communicating computers to establish a connection, or *session*, with error detection and recovery. When NetBIOS is used in datagram mode, the individual messages are sent separately in a connectionless manner, meaning that error detection and correction must be handled by the application itself.

NetBIOS also provides a name service (that is, NetBIOS names) by which computers and applications can be identified on the network.

NetBEUI is the simplest of the three protocol stacks. Its simplicity makes it the highest performer in terms of sheer speed, but the simplicity also limits its functionality. Because NetBEUI does not include a means of logical addressing for addresses at the network layer, it cannot be routed from one network or subnet to another. It works well, however, for communication within a single LAN, and it is easy to set up. It can be used in conjunction with another routable protocol such as TCP/IP. This gives you the advantages of NetBEUI's high performance within the local network and the capability to communicate beyond the LAN over TCP/IP.

IPX/SPX

IPX works in conjunction with SPX to provide routable network communications. Novell developed IPX/SPX for its NetWare servers and clients, but it can also be used with other operating systems (such as Microsoft Windows LANs). Novell based IPX/SPX on the Xerox Network System (XNS) protocols.

NOTE IPX/SPX or NWLink is required for connecting Microsoft clients to NetWare 4.x and older servers. Some NetWare 5.x servers can communicate through TCP/IP only.

Performance is higher and configuration is easier for IPX/SPX than it is for TCP/IP. IPX/SPX is sometimes used for internal LAN communications as part of a security plan. “Outside” computers accessing the LAN from the Internet, which are running only TCP/IP, are not able to access LAN systems that run only IPX/SPX.

Microsoft provides an IPX/SPX-compatible protocol stack called NWLink, which is included with all modern Windows operating systems, although it is not installed by default.

IPX works at the network layer of the OSI reference model and is connectionless. SPX operates at the transport layer of the OSI reference model and provides for acknowledgments, reassembly of packets, and other connection-oriented services.

NOTE *Connection-oriented* communication is similar to a person-to-person telephone call. If you wish to speak to Mr. Jones, you dial his number and then ask for him by name. You do not begin to communicate your message until you know you have Mr. Jones on the line (in other words, not until you have established your session).

Connectionless communication works more like using a public address system to communicate your message to a member of a large crowd. You speak into the microphone and your message is transmitted. You hope Mr. Jones was in the crowd and heard the message, but you have no way of knowing if the message did indeed reach its intended destination because no session was established.

The Network Layer Protocol: IPX


In Chapter 3, “Networking Concepts, Models, and Standards,” you learned that the network layer of the OSI reference model is responsible for logical addressing and routing

functions, that is, for getting messages to the correct destinations. This is the primary function of IPX.

For a protocol to be routable, there must be a means of identifying the network on which the computer resides. IPX uses hexadecimal *network numbers* to identify the network (subnet). A typical IPX network number would look like this: 805609a0. The administrator assigns the network number.

An IPX address consists of two parts: the network number and the *node number*, as shown in Figure 8-1. The node number identifies the specific device, and it is based on the Media Access Control (MAC) address of the interface.

Figure 8-1 An IPX address is made up of two parts: a network number and a node number.



In networks running both TCP/IP and IPX/SPX, the network numbers are often derived from the IP addresses by simply converting the IP address from decimal to hexadecimal (for example, 214.12.1.42 in hex is D6C12A).

The Role of the Service Advertising Protocol

IPX uses the Service Advertising Protocol (SAP) to advertise the addresses of network services such as file servers. A number called a SAP ID or SAP identifier is assigned to each service, and SAP broadcasts are transmitted every 60 seconds. Routers and servers keep tables that match SAP IDs to the services and update the tables dynamically with each broadcast. This keeps the tables up to date, but it also uses a great deal of network bandwidth. SAP broadcasts are not forwarded across routers, although routers can forward their SAP tables to other routers.

The Transport Layer Protocol: SPX

SPX operates above IPX, at the transport layer. Where IPX is a connectionless protocol, SPX is a connection-oriented protocol. This makes SPX more reliable, which makes sense when you think about the fact that the transport layer is responsible for acknowledgments, error checking, and other reliability issues.

IPX gets the packet to its destination. SPX concerns itself with ensuring that the packet arrives complete and in good condition. SPX handles sequencing and keeps count of the packets transmitted. It guarantees delivery by verifying the receipt of the data.

TCP/IP

The TCP/IP stack is the foundation of Internet communications. It is quickly becoming the most common network/transport solution for networks of all sizes and configurations. Thus, we focus on it in this discussion of network/transport protocols.

The following sections of this chapter talk about the concepts and theory behind the various members of the TCP/IP suite, what each one does, and how each one works. In Chapter 12, “Desktop Operating Systems,” you will learn how to configure TCP/IP for various desktop operating systems.

The TCP/IP Suite

TCP/IP is not only a protocol stack that consists of a network layer protocol and a transport layer protocol, but also a complete suite of protocols that operates at many layers of the networking model. A protocol suite, by definition, includes “extras” that are not required for network communication. These extras include, for example, the application layer utilities that are part of the TCP/IP suite.

Many of the protocols included in the suite function as information-gathering or troubleshooting utilities. In the following sections of this chapter, we examine each in turn, beginning with the main players: the network layer and transport layer protocols.

The Network Layer Protocol: IP

You know that the network layer handles routing tasks. The TCP/IP protocols enable this routing by using IP addresses to identify network devices. Every computer, network-attached printer, router, and other network device has a unique IP address.

NOTE

One device can actually have multiple IP addresses, one for each of its *network interfaces*. Routers have a minimum of two interfaces, which are attached to different subnets. Each has a unique IP address. *Multihomed* computers have more than one network interface card (NIC), and each card is assigned an IP address.

Understanding IP Addressing

Each IP address has two parts. Together, they identify the network on which the device resides and the particular device on that network. One section of the IP address represents the *network* and the second section represents the *host* (individual computer). This is much like the way in which a two-part postal address identifies a particular house to which mail is delivered:

- The street name tells the post office the general area where the house is located. Many different houses share the street name “Elm” as part of their addresses.
- The street number is unique for each individual house on that street. There can be many houses in a town with the street number “101,” but only one house on Elm Street has that number.

Similarly, many computers share the same network address, but the combined network and host address is unique to one computer (or more accurately, to one network interface). In the IP address 201.32.0.4, for example, the first three sections (called *octets*, which are discussed in the “IP Address Classes” section of this chapter) identify the network. The last section identifies the individual computer’s network interface. All computers on the same subnet have the same network ID (201.32.0), but each has a different host ID (that is, there can be only one .4 on the subnet).

If you mail a letter from Dallas to San Francisco, the post office in Dallas is not concerned with the street number in San Francisco. The first task is to get the letter to the correct city. Likewise, when you send a message across the Internet to a different LAN, the routers are not concerned with the host portion of the IP address; only the network portion is of interest. After the packet reaches the correct network (subnet), the host address is used to forward it to a specific computer—much as the local post office in San Francisco uses the street address to ensure that your letter reaches its destination.

Likewise, if a computer sends a message to IP address 201.32.0.4, in which 204.32.0 represents the network and 4 represents the host, the first step is to get the packet to the proper network. After it arrives there, it is routed internally within the network to the computer represented by the host ID (which is 4).

In our example, the first three octets identify the network. This is not always the case. In the traditional IP addressing scheme, the portion of the address that represents the network and the portion that represents the host is determined by the address *class*, as discussed in the next section.

IP Address Classes

IP addresses, like other information processed by computers, are made up of binary numbers, or *bits*. Because long strings of ones and zeros are difficult for most humans to work with, we usually denote IP addresses in *dotted-decimal* format. (See Chapter 3 if you need a review of numbering systems and how binary is converted to decimal.) The dotted-decimal

format is sometimes called *dotted quad* because there are four sets of numbers separated by dots, with each set representing an octet.

An *octet* is eight bits long; there are eight digits, each of which is a one or a zero. The four octets are sometimes designated as w.x.y.z. With this designation, the far right octet is called the “z” octet, the next one is called the “y” octet, and so on.

Because there are eight bits in each of the four octets, each IP address is a 32-bit number. This means that there are over four billion possible IP addresses (4,294,967,296, to be exact, or 2^{32}). An IP address usually looks something like this: 192.168.1.12.

NOTE

This discussion is based on version 4 of the Internet Protocol (IPv4), which is in common use today. A new proposed standard, IPv6, or IPng (which stands for IP, Next Generation), would use a 128-bit address and provide 2^{128} useable addresses.

We know that part of the IP address identifies the network and part identifies the individual device (host), but which part represents which? Unfortunately, the answer is that it depends. Traditionally, it depended upon the class to which the network belonged (a newer method of addressing, called *classless addressing*, will be discussed in the section titled “Understanding Classless Addressing”).

The Internet Assigned Numbers Authority (IANA) hands out IP addresses. In the early days of Internet communications, it seemed logical to assign IP addresses to companies and organizations in blocks because each computer on the LAN needed a unique address with which to communicate on the Internet.

NOTE

With the advent of technologies such as Network Address Translation (NAT), which was discussed in Chapter 6, “WAN Links,” it is no longer strictly true that every computer on the local network needs a public IP address.

Blocks of addresses were assigned based on the size of the local network. Large enterprise-level networks needed larger blocks of addresses, and small networks with only a few

devices required smaller blocks. *Address classes* were designated based on network size (number of host addresses). Table 8-1 shows the traditional IP address classes.

Table 8-1 *IP Address Classes*

Address Class	Number of Networks	Number of Hosts per Network
A	126*	16,777,216
B	16,384	65,535
C	2,097,152	254
D (Multicast)	N/A	N/A

* The 127.x.x.x address range is reserved as a *loopback address*, used for testing and diagnostic purposes.

NOTE

By using the entire 127.x.x.x network address for the loopback, over 24 million IP addresses were wasted. In the early days of the Internet, this was not a problem because there were far more available IP addresses than Internet computers, and the commercialization and enormous growth of the Internet was not anticipated.

As you can see in Table 8-1, there are only 126 Class A addresses available. These addresses were used up some time ago; they were assigned to very large corporations and educational institutes, including IBM, Hewlett Packard, Xerox, Massachusetts Institute of Technology (MIT), Columbia University, Digital Equipment Corporation, General Electric, and Apple. Each network has more than 16 million host addresses that can be assigned to computers within it.

The IP address class scheme creates more than 2 million Class C networks, but each of those networks can have no more than 254 host addresses. Class C network addresses are often assigned to Internet service providers, which subdivide their allocation into smaller blocks of addresses for companies that have only 10 to 20 host machines on their networks.

Class B addresses fall between Class A and Class C addresses. They are assigned primarily to large companies that were too small (or nonexistent) to have received Class A addresses back in the early days of the Internet. Microsoft Corporation is an example of a company that has a Class B network.

Class D addresses are not used for networks, but for *multicast messaging*, which is a means of sending a single message to multiple recipients simultaneously. A Class D address is assigned to a specified group of computers, and multicast protocols handle the distribution of the packets.

The method of dividing IP addresses into classes based on network size is called *classful addressing*. (The *classless addressing* method of addressing is discussed in the “Understanding Classless Addressing” section of this chapter.)

Multicasting and the MBONE

Multicast messages serve a function on the Internet that is similar to that of broadcast messages on a LAN. However, rather than being sent to *all* computers on the network, multicast messages are sent only to the computers that belong to predefined multicast groups.

There are several advantages to multicast messaging:

- It conserves bandwidth because you send the same packet to multiple addresses instead of sending the packet once for each address.
- You don't have to know the address of every computer to which you send the packet; you only have to know the multicast address.

Traditional, single-destination transmissions are called *unicast messages*. Traditional Internet routers are configured for unicast transmissions, and many cannot handle multicast packets. To solve this problem, the Internet Engineering Task Force (IETF) developed a virtual network called the MBONE (multicast backbone) that runs on top of the Internet. Using the same hardware as the Internet, the MBONE software transmits multicast packets inside unicast packets. This *encapsulation* (also referred to as *tunneling*) hides the multicast packet from routers that cannot process it and enables it to travel through the traditional nonmulticast routers.

Understanding Classful Addressing

Remember that although we commonly use decimal notation, IP addresses are actually made up of binary numbers. Table 8-2 shows how address classes can be identified based on the first octet address range. To understand Table 8-2, it is necessary to go back to the binary for a moment.

Remember that the address class is identified by the *high-order bits*, or the first few bits in the leftmost octet (sometimes referred to as the “w” octet).

Table 8-2 *Identifying Address Classes*

IP Address Class	High-Order Bits	First Octet Address Range	Number of Bits in the Network Address
Class A	0	0–127*	7
Class B	10	128–191	14
Class C	110	192–223	21
Class D	1110	224–239	28


* The 127.x.x.x address range is reserved as a *loopback address*, used for testing and diagnostic purposes.

As you can see in Table 8-2, Class A addresses are identified by the first bit of 0, Class B addresses are designated by the first two bits of 10, Class C addresses have the first three bits of 110, and multicast (Class D) addresses use 1110 as the first four bits.

Consider the following IP address: 11001111.00101100.01010001.11100111. Because the first three far-left (high-order) bits are 110, we know that this is a Class C address. If we convert the binary address to dotted-decimal notation (using a scientific calculator or the method described in Chapter 3), we get the following: 207.44.81.231.

Notice that the first octet, 207, falls into the 192–223 range shown in Table 8-2. Because all Class C addresses have 110 as the first three bits, they all have a “w” octet that falls into this range. This means that we can identify which class an IP address belongs to merely by its first octet.

In the previous example, 207.44.81 identifies the network (subnet) on which the computer resides. All computers on that subnet have these numbers as the first three octets of their IP addresses. The number 231 designates the specific host computer. No other device on that subnet can have this number as the far-right octet of its IP address. Figure 8-2 shows an example of how an IP address is divided into network and host IDs.

Figure 8-2 *In a Class C address, the first three octets represent the network ID.*

Now, how do we determine the number of networks and hosts that are available for each class? Traditionally, a Class A address uses the first octet as the network address and the remaining three as the host address. Because the very first bit (the high-order bit) in a Class A address is used to identify the address class, that leaves seven bits in the first octet that can be used for the network identification.

A Class B address uses the first two octets for the network and the last two for the host. There are 16 bits in the first two octets, but the first 2 bits are used to identify the class. Thus, we have 14 bits left for the network ID.

Following the same logic, Class C uses the first three octets (8 bits times three octets, or 24 bits) to identify the network, and only the last octet is available to identify hosts. The 3 high-order bits identify the class, so we subtract 3 from 24 and end up with 21 bits for the network ID portion of the address.

Let's go back to our Class A network. We know that 7 bits are available for the network ID. If all 7 bits are turned "on," (designated by 1s), the highest number we can have for the network ID is 1111111. When we convert this number to decimal, we have 127, which is the number of possible Class A networks (the 127 range is reserved for loopback testing). The 0.0.0.0 address is reserved for representing all IP addresses.

An easier way to arrive at this number is to raise 2 to the power of x , where x is the number of bits available for the network ID. If we use this method, we get 2^7 , or 128. When we exclude the 0.0.0.0 network address and the loopback network address, we arrive at 126, the number of possible Class A addresses.

The same process can be used for Class B and Class C networks:

- **Class B**—14 bits = $2^{14} = 16,384$
- **Class C**—21 bits = $2^{21} = 2,097,252$

NOTE

Some books tell you to use the formula $2^x - 2$. The reason for subtracting 2 is because of an old rule that said that a network ID cannot be all 0s or all 1s. This rule came into being because older routers were not able to handle network IDs of these types. Many modern routers, however, do allow for network IDs consisting of all 0s or all 1s; thus, it is not necessary to subtract 2.

The numbers in the previous bulleted list are the maximum theoretical number of networks available in each address class. Although the modern method enables *network* IDs of all 0s or all 1s, you still cannot have host IDs that consist of all 0s or all 1s.

You might have determined by now that classful addressing is not the most efficient use of the finite number of IP addresses available within the 32-bit addressing scheme. To illustrate this determination, let's use the example of a company that has 2000 computers it wants to connect to the Internet. A Class C address won't do because that would limit the network to 254 hosts. The next step up is a Class B address—but if the company obtains one, it would take over 65,000 addresses out of commission. Because it needs only 2000 of those addresses, more than 63,000 IP addresses would be wasted. *Classless addressing*, discussed in the next section, addresses this problem.

Understanding Classless Addressing

The waste of addresses associated with classful addressing has contributed to the shortage of public IP addresses. One proposed solution is the implementation of IPv6, which uses a larger address space (128 bits). However, the transition to a new version of IP is not simple and will take some time to accomplish. Meanwhile, another solution exists: classless addressing based on classless interdomain routing (CIDR).

Rather than using address classes, CIDR uses a designation appended to each IP address that specifies the number of bits used for the network portion of the address. CIDR networks are sometimes called “slash x” networks because the IP address is separated from the suffix by a slash. Thus, a CIDR address looks like this: 192.168.1.0/24. The “slash 24” means that the far-left 24 bits are used to identify the network, and the remaining eight bits are used to identify the host. In other words, the first three octets indicate the network, and the last octet specifies the host computer. In classful addressing, this would be a Class C network.

Table 8-3 shows how CIDR addresses correspond to traditional classful addresses.

Table 8-3 *CIDR (Classless) Addresses and Traditional Classful Addresses*

CIDR Address	Classful Address
/8	Class A
/16	Class B
/24	Class C

CIDR enables much more efficient allocation of IP addresses. In addition to the slash-x designations in Table 8-3, CIDR networks can be designated as /12, /20, /21, /28, and so on; that is, whatever number of bits you wish to use for the network ID can follow the slash. This enables network sizes that fall between the traditional network classes.

CIDR also supports the practice of combining small contiguous blocks of network addresses into one larger one. This is called *supernetting* and is discussed later in this chapter in the section “IP Subnetting and Supernetting.”

Automatic Address Allocation

To communicate using TCP/IP, a computer or other network device must have a unique IP address. This is a *logical* address and is processed at the network layer.

The network portion of the address must be the same as that of other computers on its subnet. For example, if you were using the default subnet mask for Class C networks, 192.168.1.12 and 192.168.1.34 would be two computers on the same subnet because the network ID, represented by the first three octets, is the same.


In contrast, the host portion must *not* be the same as that of any other computer on the same subnet. For example, there could not be two computers with host address .6 on the same subnet.

There are two ways to obtain an IP address:

- The address can be manually entered into the operating system’s TCP/IP properties configuration. This requires that the network administrator assigning the address understand TCP/IP addressing and know how to choose a valid address for the particular network.
- The address can be automatically assigned. Generally, this means a computer on the network is configured as a DHCP server to hand out IP addresses from a pool of valid addresses. In other cases, an operating system feature called Automatic Private IP Addressing (APIPA) enables a computer to assign itself an address if it is unable to contact a DHCP server.

DHCP

DHCP is a protocol that runs on a machine designated as a DHCP server, which allocates IP addresses to machines that are configured to be DHCP clients. Figure 8-3 illustrates the process.

Figure 8-3 *The DHCP server leases an address to the DHCP client.*

The steps involved when a DHCP client obtains an IP address lease from a DHCP server is as follows:

- 1 A computer whose TCP/IP properties are set to obtain an IP address through DHCP comes onto the network. This computer broadcasts a message called a DHCP discover message to the entire network or subnet. (Broadcast messages are sent to all computers by using a special broadcast address.)
- 2 If a DHCP server is present on the network, it receives the broadcast message and responds with a message called a DHCP offer. The message offers the client an IP address from the server's range of addresses that are available for allocation. The offered address is temporarily reserved until the server receives a response from the client. The offer message is also sent as a broadcast because the client does not yet have an IP address of its own to which a one-to-one message can be sent.
- 3 The client can receive multiple offers if there are multiple DHCP servers on the network. When the first offer arrives, the client responds with a message called a *DHCP request*. This is an acceptance of the first offer received. Again, it is a broadcast message; thus, all DHCP servers receive it and those who made late offers will subsequently know that their offers were not accepted. They can then place the offered addresses they had reserved for the client back into their available address pools.
- 4 The last step in this negotiation is the *DHCP acknowledgment* (also called the *ACK*). The DHCP server whose offer was accepted receives the client's request message. The server acknowledges the acceptance and assigns the IP address to that client for the

duration of a preset “lease” period. It can send the client additional TCP/IP configuration information, such as the IP addresses of DNS and WINS servers. (We discuss the role of those servers later in this chapter.)

When these steps are completed, the client can use the assigned IP address to communicate with other computers running the TCP/IP protocols and do so until the lease period expires. Note that the length of the lease can be set by an administrator on the DHCP server.

Before the expiration of the lease, the client begins negotiations to renew it so that it can continue to use the address. Normally the DHCP server grants this request. If the DHCP server has gone offline, however, or if the server sends a Negative Acknowledgment response (NACK), the client must start the DHCP process all over again.

NOTE

The DHCP server might send a NACK if, for example, the client machine has moved to a different subnet and the address it is trying to renew is no longer valid for its location.

DHCP has many advantages over manual IP addressing:

- It saves time because the administrator does not have to enter the addresses into each computer’s property settings.
- It ensures greater accuracy because the administrator does not have to keep up with which addresses have already been assigned and which are still free.

If computers on a network must have the same IP address (this is often true of servers), they can still use DHCP. In this case, you must configure those computers to use a *reserved address*. The DHCP server always assigns the same address to a client that has such a reservation. The reservation is made based on the DHCP client’s MAC (that is, physical) address.

The Origins of DHCP

DHCP grew out of an earlier protocol, the Bootstrap Protocol (BOOTP). BOOTP was originally developed to enable diskless workstations to boot up, be assigned an IP address, and then load an operating system over the network.

DHCP is much more advanced than BOOTP, and it enables the administrator to configure many options and set lease durations. DHCP adds dynamic allocation of addresses. Some DHCP servers are configured to support BOOTP clients.

DHCP is not operating system specific. It can be used with Microsoft, UNIX, NetWare, and other popular network types. However, vendor implementation of the DHCP services may

differ. For example, Windows 2000 DHCP servers are integrated with Active Directory. This enables administrators to prevent unauthorized DHCP servers (sometimes referred to as “rogue” DHCP servers) from handing out IP addresses on the network.

APIPA

Another means of automatically obtaining an IP address is APIPA. The TCP/IP implementation of recent Microsoft operating systems, such as Windows 98 and Windows 2000, include this feature.

Traditionally, if a computer was configured to be a DHCP client, and it was unable to contact a DHCP server when it came onto the network, that computer would have no IP address and would not be able to communicate over TCP/IP. APIPA was introduced to solve this problem.

When an APIPA-enabled computer cannot locate a DHCP server to obtain an address, it assigns itself one from a range of addresses reserved for that purpose (the Class B 169.254.0.0 network range). The self-assigned address can be used until the DHCP server is functional again.

IP Subnetting and Supernetting

To *subnet* a network means to divide it into parts. Subnetting turns the two-level address hierarchy described earlier into a three-level addressing system. To subnet a network, you “borrow” some of the bits that are normally used for the host portion of the address and use them for the second level of the network address (that is, the subnet address).

The Subnet Mask

Subnetting involves borrowing from one part of the address to give to the other. When we subnet, IP must have a means of determining which bits identify the network and which are still being used to indicate the host. IP determines this by using the *subnet mask*, a 32-bit number entered by the administrator in the TCP/IP properties configuration. The subnet mask sets the bits representing the network ID to 1s and those representing the host ID to 0s. The network ID portion of the address is said to be “masked” by the bits that are turned on.

By default, Class A networks use the bits in the first octet for the network ID, Class B networks use the bits in the first two octets to identify the network, and Class C addresses

use the first three octets for this purpose. The *default subnet masks* resulting from this pattern are shown in Table 8-4.

Table 8-4 *Default Subnet Masks*

Address Class	Binary Subnet Mask	Decimal Subnet Mask
Class A	11111111.00000000.00000000.00000000	255.0.0.0
Class B	11111111.11111111.00000000.00000000	255.255.0.0
Class C	11111111.11111111.11111111.00000000	255.255.255.0

The subnet masks in Table 8-4 apply to *unsubnetted* networks, and as long as we stick with the defaults, this is simple enough.

What happens, however, if we want to divide a network? Let's say we have been assigned a Class B network address, such as 181.25.0.0. We know that a Class B network can contain 65,535 host computers. However, if we had that many computers on one network, broadcast traffic would be unmanageable.

As a solution, let's suppose we have decided to divide the network into six subnets. To do so, we must borrow bits from the host address portion of the address, which is used to indicate the subnets, and we must calculate the correct subnet mask that indicates to IP that our network has six subnets (or eight possible subnets).

Calculating the Subnet Mask

A subnet mask other than the default masks is referred to as a *variable-length* or *custom* subnet mask. To calculate the correct subnet mask for our scenario in which six separate subnets are desired, we must first determine how many bits we need to borrow from the host portion of the address.

Because binary is a base 2 numbering system, subnets must be created in blocks of powers of 2. To calculate the subnet mask, we must find out what power of 2 gives us 6 (or more) subnets.

If we raise 2 to the second power (2×2), we get 4. We need more subnets than that, so let's try 2 to the third power ($2 \times 2 \times 2 = 8$). If we subtract 2, to comply with the old rule that says we can't use network IDs that consist of all 0s or all 1s, we have 6 remaining usable subnets. Therefore, we need to borrow three bits from the host portion of the address. This means we must turn the first three 0s, which indicate the host ID, into 1s, which indicate the network ID. Our original default subnet mask, with the borrowed bits, now looks like this: 11111111.11111111.11100000.00000000. If we convert it to decimal, we have this: 255.255.224.0.

How many host computers can we have on each subnet? Look at the remaining 0s that indicate the host portion of the address. You'll find there are 13 of them, which gives us $2^{13} = 8192$. This means that we can have 8190 hosts on each subnet, after subtracting 2 so that no host address is all 0s or all 1s. (Remember that the prohibition against all 0s and 1s has been removed only in regard to network IDs. The host IDs still follow this rule.)

Table 8-5 provides a quick reference for the number of subnets and hosts enabled with each subnet mask.

Table 8-5 Quick Reference Subnetting Chart

Decimal Notation for First Octet	Number of Subnets	Number of Class A Hosts	Number of Class B Hosts	Number of Class C Hosts
.192	2	4,194,302	16,382	62
.224	6	2,097,150	8190	30
.240	14	1,048,574	4094	14
.248	30	524,286	2046	6
.252	62	262,142	1022	2
.254	126	131,070	510	—
.255	254	65,534	254	—

ANDing

When a network transmission is sent using TCP/IP, IP must determine whether the destination computer is on the same subnet as the sending computer. If both computers are on the same subnet, the message is broadcast. If the destination computer is on a different subnet, the message is sent to the *default gateway* address, which is the address of the router's interface (the router serves as the gateway out of the subnet).

IP uses a process called ANDing to ascertain whether the sending and destination computers are on the same subnet. ANDing is done by combining the binary versions of the IP address for each computer with the subnet mask. In combining these binary numbers, the calculations are made as follows:

- 1 AND 1 = 1
- 1 AND 0 = 0
- 0 AND 0 = 0

Here is an example of ANDing: The IP address for the sending computer is 192.168.1.1, with a subnet mask of 255.255.255.0. The IP address for the destination computer is 192.168.3.1 with a subnet mask of 255.255.255.0.

First, we AND the sending computer's IP address with the subnet mask:

```
192.168.1.1 = 11000000.10101000.00000001.00000001 255.255.255.0 =
11111111.11111111.11111111.00000000
ANDed result =11000000.10101000.00000001.00000000
```

Then, we do the same calculation for the destination computer and subnet mask:

```
192.168.3.1 = 11000000.10101000.00000011.00000001 255.255.255.0 =
11111111.11111111.11111111.00000000
ANDed result =11000000.10101000.00000011.00000000
```

The results are different, so we know (and more importantly, IP knows) that these two computers reside on different subnets. The message is then sent to the router (default gateway) to be forwarded to the correct subnet.

If the result had been the same, IP would know the destination computer was on the local subnet and the message would be sent using an Address Resolution Protocol (ARP) broadcast. We discuss more about ARP later in this chapter.

Benefits of Subnetting

The benefits of separating a large network into two or more subnetworks (subnets) include the following:

- It reduces broadcast traffic. Subnets are connected to one another by routers, and most routers are configured by default not to pass on broadcast messages. This can substantially conserve network bandwidth.
- It organizes computers at different locations into separate subnets for easier management.
- It isolates a part of the network for security or filtering purposes.
- It provides more efficient use of available addresses and fewer "wasted" addresses.

IP subnetting is the subject of entire books. It is a complex topic and all the details involved are beyond the scope of this introductory text (the list at the end of this chapter has several good subnetting tutorials). This section provided, however, a very basic overview of subnetting concepts.

The Transport Layer Protocols: TCP and UDP

Recall from Chapter 3 that the transport layer (called the *host-to-host* layer in the DoD model) is responsible for providing reliable end-to-end communication. This is accomplished by mechanisms such as acknowledgments, which verify that data has arrived at the destination without damage or loss.

Transport layer protocols also differentiate between messages that arrive at the same destination computer. Because different applications can be sending or receiving messages

at the same time, the transport layer protocols use *ports* to keep these messages separate. We discuss ports and a related concept, *sockets*, later in this section.

The TCP/IP suite includes not one, but two transport layer protocols:

- **Transmission Control Protocol (TCP)**—A connection-oriented protocol
- **User Datagram Protocol (UDP)**—A connectionless protocol

Which of the two transport protocols is used to send a particular message? It depends on the needs for that transmission. TCP is appropriate when reliability is of utmost importance, and UDP is used if performance (speed) is the highest priority. We examine the characteristics and functions of each in the following sections.

The Transport Layer Protocol: TCP

TCP, because it is a connection-oriented protocol, establishes a session between the two communicating computers before sending data. Acknowledgment and response messages are used to establish the session. Error checking and correction are then performed, and the data is broken down into packets.

Sequencing information is added to each packet so that the parts of the message can be put back together in the correct order. This information also enables the receiving computer to detect if packets are missing. This makes TCP more reliable than UDP, but at a price: all these extra duties slow down performance.

The Transport Layer Protocol: UDP

UDP is connectionless. It does not sequence the packets in which data arrives; this means it is more appropriate for small messages that can be transmitted in one packet. UDP also doesn't keep track of what it has sent or guaranteed. It does provide for a checksum, however, to ensure that data is intact upon arrival. Like TCP, it provides port numbers to differentiate between the requests sent by or being delivered to different applications.

Because it does not have to bother with sequencing and error checking, UDP is fast. Its header is less complex than the header added by TCP. The Routing Information Protocol (RIP), the Trivial File Transfer Protocol (TFTP), and name lookup messages use UDP.

NOTE

What is a *datagram*? The term is defined in RFC 1594 as “a self contained, independent entity of data carrying sufficient information to be routed from the source to the destination computer without reliance on earlier exchanges between this source and destination computer and the transporting network.”

The terms datagram and *packet* are sometimes used interchangeably. “Packet” describes a unit of data that is part of a sequenced group of units or “chunks” into which a message is broken. Packets can take different routes over the network to reach the destination. They are then reassembled at the destination. Datagram is used to describe the simpler, nonsequenced data units transmitted by UDP.

Ports and Sockets

TCP/IP uses a two-part logical address—the IP address—to identify the source and destination computers in network communication. What happens, however, if two network applications running on the same computer are sending requests and receiving responses simultaneously? For example, what happens if one incoming message is intended for your e-mail program, while another is a Web page being returned to your Web browser? The protocols need a way to differentiate between them. That’s where TCP and UDP ports come in.

The Role of Port Numbers

Remember that the IP address of a destination computer contains two parts: a network address that functions somewhat like a street name, and a host address that functions somewhat like a street number. You can think of port numbers as specific routing information within an address. It is an addendum to the IP address, just as the name of an addressee on an envelope is an addendum to the street address. Likewise, you can think of separate applications as separate residents.

A port is a *logical connection point*. Ports are used by the transport protocols, TCP and UDP, to identify the specific application that is sending or receiving the message.

Commonly used Internet applications have predefined port numbers. This standardization makes communication easier. The preassigned port numbers are called *well-known ports*. Table 8-6 lists examples of these commonly used port numbers.

Table 8-6 TCP/UDP Well-Known Ports

Preassigned Port	Protocol	Application
80	TCP	HTTP
21	TCP/UDP	FTP
23	TCP/UDP	Telnet
25	TCP/UDP	SMTP
110	TCP/UDP	POP3
119	TCP/UDP	NNTP

continues

Table 8-6 *TCP/UDP Well-Known Ports (Continued)*

137	TCP/UDP	NetBIOS name service
161	TCP/UDP	SNMP
194	TCP/UDP	IRC
389	TCP/UDP	LDAP
396	TCP/UDP	NetWare over IP
458	TCP/UDP	Apple QuickTime
500	TCP/UDP	ISAKMP

There are 65,536 useable ports. Ports 0 through 1024 (the “well-known ports”) are reserved for predefined services, such as the ones shown in Table 8-6.

What Is a Socket?

Now that you understand the function of ports, we will discuss the concept of sockets. The common definition of socket is “the endpoint of a connection”; a socket must be created for communication to take place.

Different socket types use different addressing methods. The most common method uses an IP address combined with a port number to identify the socket. In UNIX terminology, this is called `AF_INET` addressing. A second addressing method, `AF_UNIX`, uses pathnames to identify the sockets.

Berkeley (BSD) Sockets became the standard API for TCP/IP communications. A popular adaptation of the sockets interface is *Windows Sockets*, or *Winsock*. This implementation provides an API for Internet applications running on Windows operating systems. Winsock is loaded as a dynamic link library (DLL).

Addressing the Envelope: Packet Headers


You now know that the TCP/IP protocols use addressing information to get a message to the correct destination. You might be wondering, however, where the protocols find all this addressing information. Just as a letter you send through the postal service must be placed in an envelope and the sender’s and recipient’s addresses noted on the outside, a message sent through TCP/IP must be placed “inside an envelope.”

When the message is placed in the “envelope,” the data is encapsulated in protocol *headers*, which contain the addresses. The headers can also include other information and special instructions, just as you can write “Handle with Care” or “Special Delivery” on your postal envelope to aid in proper delivery.

By default, the IP header is 20 bytes long, and it includes fields that indicate the following (see Figure 8-4 for more details):

- Type of service
- Total length of the datagram
- Unique identification of the datagram
- Flags and fragmentation offset to aid in reassembly
- Time To Live (TTL) to limit the number of routers through which the datagram can pass
- Upper-layer protocol that is to receive the data (ICMP, TCP, UDP, IGRP, or OSPF)
- Checksum for detection of corruption
- Source IP address (sending computer)
- Destination IP address (receiving computer)

Figure 8-4 The IP header is made up of 12 fields, plus options.



The header can include options such as security restrictions, timestamps, and routing restrictions. The header *without* options is 20 bytes long.

Name Resolution

We have discussed how TCP/IP uses IP addresses and port numbers to identify networks, computers, and specific network applications to which messages are sent. Most human beings, however, prefer to use names instead of numbers for identification purposes. That

preference is why we never want to imagine having to memorize our friends' social security numbers to designate to whom we are talking!

Although we like to use names when we access a computer on the network or type a Web server location into a Web browser, computers can work only with numbers. Because of this incompatibility, we need services that translate “friendly” names into IP addresses. Using these services, we can type `www.xerox.com` into our browser's address box instead of `208.134.240.50` when we want to view the Xerox Web site. It's certainly easier to remember. Regardless of our actions, however, our browser is converting the host name to an IP address to find the Web server on the Internet and to retrieve the requested page.

What's in a Name?

Different name types are used in computer network communications. The Internet arranges host (computer) names in a hierarchical structure within *domains*. The most common top-level domains in the United States are as follows:

- **com**—Originally intended for commercial organizations
- **net**—Originally intended for networks such as ISPs
- **org**—Originally intended for nonprofit organizations
- **edu**—Restricted for use by educational institutions
- **gov**—Restricted for use by U.S. governmental entities
- **mil**—Restricted for use by U.S. military units
- **int**—Restricted for use by international organizations

Outside the United States, the following country codes are used to identify domains:

- **uk**—United Kingdom
- **au**—Australia
- **ca**—Canada

Businesses, organizations, and individuals register *second-level domain names*, for example, `ibm.com`, `whitehouse.gov`, or `dallas.net`, within these top-level domains. At one time, second-level domain names were assigned by InterNIC, but that task has now been distributed to several authorized name registrars.

Within your second-level domain, individual computers are identified by the host name, the second-level domain name, and the top-level domain name—with dots separating each section. Thus, a Web server named “www” in the `dallas.net` domain is identified as `www.dallas.net`. This “dotted” hierarchical name is called the *fully qualified domain name (FQDN)*.

On Microsoft networks, each computer is also assigned a NetBIOS name. This is a 16-character name, assigned by the administrator, used to identify resources on the local network.

Both types of names must be translated into IP addresses before TCP/IP communication can take place.

Translating Names to Numbers

You can use the following to translate names into IP addresses:

- HOSTS and LMHOSTS files, which are text files stored on computer hard disks
- DNS or Dynamic DNS (DDNS)
- WINS

The following sections describe these means of translation.

HOSTS and LMHOSTS Files

In the early days of the Internet, the method for matching host (computer) names to IP addresses for TCP/IP communication was the HOSTS file. This is a text file, stored on the local hard drive, which lists host names and their corresponding IP addresses. Example 8-1 shows an example of a HOSTS file.

Example 8-1 *A Local HOSTS File Maps IP Addresses to Host Names*

```
102.54.94.97  rhino.acme.com  # source server
38.25.63.10   x.acme.com      # x client host
127.0.0.1     localhost
```

NOTE

The pound sign (#) in the HOSTS file indicates a comment. Any information that follows the # will be ignored by the computer.

The HOSTS file worked when there were only a few computers on the Internet. It is simple to construct and it can be modified with any text editor (for example, Windows Notepad, and vi or Emacs in UNIX or Linux). When a user attempted to access another computer using its “friendly” host name, the operating system consulted the HOSTS file for the “real” identification, the IP address.

However, each time another host was added to the network, the HOSTS file had to be manually updated, and the new file saved to all the computers on the network. As the Internet grew, this became an impossible task.

LMHOSTS serves a similar function in Microsoft networks by mapping IP addresses to NetBIOS names instead of to computer names. It suffers from the same disadvantages as HOSTS, however; it is a static file, and it must be manually updated. See Example 8-2 for an illustration of a typical LMHOSTS file.

Example 8-2 *The LMHOSTS File Matches IP Addresses to NetBIOS Names*

```
102.54.94.97 rhino #PRE #DOM:networking #net group's DC
102.54.94.102 "appname \0x14" #special app server
102.54.94.123 popular #PRE #source server
102.54.94.117 localsrv #PRE #needed for the include
```

A better means of matching up names to IP addresses was obviously needed.

DNS and DDNS

The DNS was devised to solve the problems inherent in using HOSTS files. DNS servers store databases of IP-to-hostname mappings, and clients' TCP/IP properties are configured with the address of the DNS server. When a friendly hostname needs to be translated to its IP address, the client contacts the DNS server.

A hierarchy of DNS servers exists on the Internet, with different servers maintaining DNS information for their own "zones," or areas of authority. If the DNS server consulted by your computer does not have an IP mapping for the hostname you entered, it can pass the query to another DNS server until the information is obtained.

DNS is not absolutely required to communicate on the Internet, but without it, all communications must use IP addresses instead of hostnames. For example, if you do not have a DNS server address configured in your computer's TCP/IP properties, you can still access a Web site by typing its IP address into the URL field. However, if you type in the hostname instead, the browser is unable to return the page. The DNS server address can be entered manually, or it can be obtained from a DHCP server if your computer is set up as a DHCP client.

DNS is a big improvement over local HOSTS files because the database is stored on a central server and you need only update it there instead of on all client machines. However, the server's database still must be updated manually. *Dynamic DNS* addresses this problem by enabling automatic updates of the DNS database. Using this enhanced form of DNS, client computers can register and update their resource records on the DNS server when changes occur.

Windows 2000 DNS servers support the DDNS protocol extension, as does BIND version 8. RFC 2136 contains specifications for DDNS standards.

The DNS Database Table

DNS uses different types of records in the database table. The following are some of the common record types:

- **Address (A) Record**—Maps a host name to an IP address
- **Mail Exchange (MX) Record**—Points to a mail exchange server for a specific host
- **Canonical Name (CNAME) Record**—Maps *aliases*, or additional names, to a host

All these record types (and others for specialized purposes) are combined in the DNS table.

WINS

WINS is another method for resolving names to IP addresses. In this case, NetBIOS names (used to identify computers and services on Microsoft networks) are mapped in a database on a WINS server. Windows NT and Windows 2000 servers can function as WINS servers.

NetBIOS names are flat rather than hierarchical as are FQDNs (DNS host names). Where the FQDN for a particular server might be *exeter.tacteam.net*, a NetBIOS name would be simply *Exeter*. TCP/IP doesn't understand NetBIOS names; again, it needs an IP address to communicate with the server. WINS, like the static LMHOSTS file, can translate the name to the required IP number.

WINS, unlike the original DNS, uses a dynamically updated database. When WINS clients come onto the network, they announce themselves to the WINS server, giving their names and IP addresses. The WINS server builds its database from this information.

DHCP, DNS, and WINS can all work together on the same network. In new operating systems such as Microsoft Windows 2000, the three services are integrated to interoperate efficiently.

Summarizing Name Resolution Methods

It is easy to confuse the various name resolution methods. Table 8-7 summarizes the features and uses of each.

Table 8-7 *Name Resolution Methods*

Name Resolution Method	Name Type Resolved	Characteristics
HOSTS file	Host names to IP addresses	Text file; must be updated manually on each computer
LMHOSTS file	NetBIOS names to IP addresses	Text file; must be updated manually on each computer

continues

Table 8-7 *Name Resolution Methods (Continued)*

DNS	Host names to IP addresses	Centralized database managed by DNS server; must be updated manually
DDNS	Host names to IP addresses	Centralized database managed by DNS server; can be updated dynamically
WINS	NetBIOS names to IP addresses	Centralized database managed by WINS server; can be updated dynamically

TCP/IP Utilities

TCP/IP is a complex collection of protocols. Most vendors' implementations of the suite include a variety of utilities for viewing configuration information and troubleshooting problems. In the following section, we look at the following common TCP/IP utilities:

- Packet Internet groper (ping)
- Address Resolution Protocol (ARP) and Reverse ARP (RARP)
- Netstat and tpcon
- Nbtstat
- IP configuration utilities: ipconfig, winipcfg, config, and ifconfig
- Route-tracing utilities: traceroute, tracert, and iptrace

NOTE Utilities that perform the same function(s) may be given different names by different vendors.

Ping

A simple but highly useful command-line utility included in most implementations of TCP/IP is ping. Ping can be used with either the hostname or the IP address to test IP connectivity.

Ping works by sending an ICMP echo request to the destination computer. That receiving computer then sends back an ICMP echo reply message.

It is also possible to use ping to find out the IP address of a host when you know the name. If you type the **ping apple.com** command as shown in Example 8-3, you will see the IP address from which the reply is returned.

Example 8-3 *Information Returned in Response to a ping of apple.com*

```
c:\>ping apple.com
Pinging apple.com [17.254.3.183] with 32 bytes of data:
Reply from 17.254.3.183: bytes=32 time=430ms TTL=90
Reply from 17.254.3.183: bytes=32 time=371ms TTL=90
Reply from 17.254.3.183: bytes=32 time=370ms TTL=90
Reply from 17.254.3.183: bytes=32 time=371ms TTL=90

Ping statistics for 17.254.3.183:
    Packets: Sent = 4, Received = 4, Lost = 0 <0% loss.,
    Approximate round trip times in milli-seconds:
        Minimum = 370ms, Maximum = 430ms, Average = 385ms
```

Another utility, nslookup, returns the IP address for a given host name and a host name for a given IP address.

Novell implements ping as a NetWare Loadable Module (NLM). Windows and UNIX/Linux operating systems use the **ping** command at the command line. Third-party ping utilities are available, some of which provide a graphical interface.

ARP and RARP

ARP refers to the protocol itself and to the command-line utility used to view and manipulate the ARP cache. You must understand the function of the protocol to properly use the utility.

ARP: The Protocol

ARP is the means by which networked computers map logical IP addresses to physical hardware (MAC) addresses. ARP builds and maintains a table called the *ARP cache*, which contains these mappings. RARP is used by a machine that doesn't know its IP address to obtain the information based on its MAC address.

ARP: The Utility

ARP is also a command-line utility provided with Windows and UNIX/Linux TCP/IP stacks that can be used to view and change ARP's IP-to-MAC address mappings. (Novell

implements this as an NLM called `tpcon`.) With the ARP utility, you can display the contents of the cache and add or delete specific mappings, as shown in Example 8-4.

Example 8-4 *The arp Command Is Used to View the ARP Cache*

```
c:>arp -a
Interface: 192.168.1.201 on Interface 0x2
  Internet Address      Physical Address      Type
  192.168.1.16         00-40-f6-54-d7-43    dynamic
  192.168.1.185        00-50-da-0d-f5-2d    dynamic
```

The following switches can be used with the `arp` command:

- `arp -a`—Displays the cache
- `arp -s`—Adds a permanent IP-to-MAC mapping
- `arp -d`—Deletes an entry

There are other switches included with specific vendors' implementations of ARP.

Netstat/Tpcon

It is often useful to view network statistics. The `netstat` command is used in Windows and UNIX/Linux to display TCP/IP connection and protocol information. Novell uses the `tpcon` NLM to accomplish this.

The `netstat` command provides a list of connections that are currently active, as shown in Example 8-5.

Example 8-5 *The netstat Command Is Used to View Connection Information*

```
c:>netstat
Active Connections
 Proto Local Address      Foreign Address      State
 TCP   DS2000:3301        msgr-ns18.hotmail.com:1863 ESTABLISHED
 TCP   DS2000:3450        constellation.tacteam.net:3389 ESTABLISHED
 TCP   DS2000:3860        ultra1.dallas.net:pop3 TIME_WAIT
 TCP   DS2000:3861        aux153.plano.net:pop3 TIME_WAIT
```

In Example 8-5, you can see the protocol used for each connection, the local computer name and port number used for the connection, the “foreign” address (the remote computer name), and the state of the connection.

Several switches can be used with `netstat`, as shown in Table 8-8.

Table 8-8 *netstat Switches and Their Functions*

Switch	Function
-a	Shows all connections and listening ports
-e	Shows Ethernet statistics

Table 8-8 *netstat Switches and Their Functions*

-n	Shows addresses and ports
-p*	Enables you to display information only for selected protocol
-t, -u, -w, -x†	Enables you to display information for TCP, UDP, RAW, or sockets
-r	Shows the routing table
-s	Provides a summary of statistics for each protocol

* Used with Microsoft TCP/IP implementation

† Used with Linux TCP/IP implementation

Netstat statistics can be useful in troubleshooting TCP/IP connectivity problems. Example 8-6 shows the wealth of information available in summary (-s switch) mode. These error reports are especially helpful in diagnosing hardware and routing problems.

Example 8-6 *The netstat -s Command Displays TCP/IP Statistics*

```

C:>netstat -s
IP Statistics
Packets Received                = 1091043
Received Header Errors          = 0
Received Address Errors         = 7
Datagrams Forwarded             = 0
Unknown Protocols Received      = 0
Received Packets Discarded      = 0
Received Packets Delivered      = 1091034
Output Requests                 = 420049
Routing Discards                = 0
Discarded Output Packets        = 0
Output Packets No Route         = 0
Reassembly Required             = 6
Reassembly Successful           = 3
Reassembly Failures             = 0
Datagrams Successfully Fragmented = 12
Datagrams Failing Fragmentation = 0
Fragments Created               = 24
ICMP Statistics
                                     Received   Sent
Messages                            994       1129
Errors                               0          0
Destination Unreachable             12         82
Time Exceeded                       0          0
Parameter Problems                  0          0
Source Quenches                     0          0
Redirects                            0          0
Echoes                              37       1010
Echo Replies                        945         37
Timestamps                          0          0
Timestamp Replies                   0          0

```

continued

Example 8-6 *The netstat -s Command Displays TCP/IP Statistics (Continued)*

Address Masks	0	0
Address Mask Replies	0	0
TCP Statistics		
Active Opens	=	3940
Passive Opens	=	42
Failed Connection Attempts	=	77
Reset Connections	=	930
Current Connections	=	2
Segments Received	=	577343
Segments Sent	=	388999
Segments Retransmitted	=	361
UDP Statistics		
Datagrams Received	=	38481
No Ports	=	475173
Receive Errors	=	0
Datagrams Sent	=	29404

Nbtstat

The Microsoft TCP/IP stacks included in Windows operating systems provide the `nbtstat` utility, which is used to display NetBIOS information. Example 8-7 shows the syntax and switches available with the `nbtstat` command.

Example 8-7 *Type nbtstat at the Command Line for a Display of the Syntax and a List of Available Switches*

```
c:\>nbtstat

Displays protocol statistics and current TCP/IP connections using NBT
(NetBIOS over TCP/IP).

NBTSTAT [ [-a RemoteName] [-A IP address] [-c] [-n]
          [-r] [-R] [-RR] [-s] [-S] [interval] ]

-a (adapter status) Lists the remote machine's name table given its name
-A (Adapter status) Lists the remote machine's name table given its
IP address.
-c (cache)          Lists NBT's cache of remote [machine] names and their IP
addresses
-n (names)          Lists local NetBIOS names.
-r (resolved)       Lists names resolved by broadcast and via WINS
-R (Reload)         Purges and reloads the remote cache name table
-S (Sessions)       Lists sessions table with the destination IP addresses
-s (sessions)       Lists sessions table converting destination IP
addresses to computer NETBIOS names.
-RR (ReleaseRefresh) Sends Name Release packets to WINS and then, starts
Refresh

RemoteName  Remote host machine name.
```

Example 8-7 *Type nbtstat at the Command Line for a Display of the Syntax and a List of Available Switches (Continued)*

IP address	Dotted decimal representation of the IP address.
interval	Redisplays selected statistics, pausing interval seconds between each display. Press Ctrl+C to stop redisplaying statistics.

Ipconfig, Winipcfg, Config, and Ifconfig

TCP/IP configuration information can be displayed using the following utilities, depending on the operating system:

- **Ipconfig**—Windows NT and Windows 2000 (command-line)
- **Winipcfg**—Windows 95 and 98 (graphical interface)
- **Ifconfig**—UNIX and Linux (command-line)
- **Config**—NetWare (server console)

The configuration utilities can provide a wealth of information, including currently used IP address, MAC address, subnet mask, and default gateway; addresses of DNS and WINS servers; DHCP information; and services enabled. There is a variety of switches available, depending on the vendor and specific utility. See Example 8-8 for the results of using the **ipconfig** command with the **/all** switch in Windows 2000.

Example 8-8 *Configuration Information Is Displayed by the Windows 2000 ipconfig /all Command*

```
c:\>ipconfig/all
Windows 2000 IP Configuration

Host Name . . . . . : DS2000
Primary DNS Suffix . . . . . : tacteam.net
Node Type . . . . . : Hybrid
IP Routing Enabled. . . . . : Yes
WINS Proxy Enabled. . . . . : No
DNS Suffix Search List. . . . . : tacteam.net

Ethernet adapter Local Area Connection:

Connection-specific DNS Suffix . :
Description . . . . . : 3Com EtherLink XL 10/100 PCI TX NIC
(3C905B-TX)
Physical Address. . . . . : 00-50-04-7C-C0-D2
DHCP Enabled. . . . . : No
IP Address. . . . . : 192.168.1.201
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . : 192.168.1.16
DNS Servers . . . . . : 192.150.87.2
                          216.87.128.131
Primary WINS Server . . . . . : 192.168.1.185
```

Tracert, Iptrace, and Traceroute

It is often useful to trace the route a packet takes on its journey from source computer to destination host. TCP/IP stacks include a route tracing utility that enables you to identify the routers through which the message passes. Depending on your operating system, you can use one of the following:

- **Tracert**—Windows
- **Iptrace**—NetWare NLM
- **Traceroute**—UNIX/Linux

Example 8-9 show the results of a trace using the **tracert** command, which has the following syntax:

```
tracert destination hostname
```

Example 8-9 You Use tracert to Trace the Route of a Packet from Source to Destination

```
c:\>tracert dallas.net

Tracing route to dallas.net [204.215.60.1]
over a maximum of 30 hops:

  1  <10 ms  <10 ms  <10 ms  STARBLAZER [192.168.1.16]
  2  60 ms   90 ms   60 ms   dal-isdn0.august.net [216.87.128.117]
  3  80 ms   60 ms   100 ms  dal-gw-eth0.august.net [216.87.128.126]
  4  60 ms   70 ms   61 ms   dal-gw2-fe2-0.august.net [216.87.128.86]
  5  60 ms   70 ms   100 ms  500.Serial2-7.GW6.DFW9.ALTER.NET [157.130.216.157]
  6  80 ms   70 ms   60 ms   158.at-5-0-0.XR1.DFW9.ALTER.NET [152.63.100.186]
  7  60 ms   110 ms  70 ms   185.ATM6-0.XR1.DFW4.ALTER.NET [152.63.96.137]
  8  80 ms   70 ms   80 ms   195.ATM11-0-0.GW1.DFW1.ALTER.NET [146.188.240.41]
  9  90 ms   320 ms  171 ms  savvis-dfw-gw.customer.ALTER.NET [157.130.128.54]
 10  90 ms   90 ms   100 ms  ETHOS-1.usd11s.savvis.net [209.44.32.10]
 11 150 ms  130 ms  131 ms  cisco-plano-e0.dallas.net [204.215.60.1]

Trace complete.
```

As you can see, the trace shows the IP address and the name of the forwarding computer or router. The packet required five hops to reach its destination, which was a host named `www.dallas.net`. Roundtrip times (in milliseconds) are shown for each hop.

Summarizing TCP/IP Utilities

TCP/IP is a large and complex suite of protocols. Most implementations include a variety of utilities that can be used for information gathering and troubleshooting. Table 8-9 offers a summarization of the common utilities.

Table 8-9 *TCP/IP Utilities*

Utility	Use
ARP/RARP	To view IP address to MAC address entries that have been resolved by ARP protocol, to delete entries from the ARP cache, and to add permanent IP-to-MAC mappings
Netstat (Windows/UNIX), tpcon (NetWare)	To view network connections and protocol statistics
Netbtstat (Windows)	To view connections and statistics for NetBIOS over TCP/IP (NetBT)
Ipconfig (Windows NT/2000), Winipcfg (Windows 95/98), Config (NetWare), Ifconfig (UNIX)	To view TCP/IP configuration information such as IP address, subnet mask, default gateway, MAC address, services enabled, and more
Tracert (Windows), Iptrace (NetWare), Traceroute (UNIX)	To discover the route taken by a packet on its journey from the source to the destination computer and to identify the routers through which it passes
Ping	To determine IP connectivity between two systems

Application Layer Protocols

The TCP/IP suite includes a variety of application layer protocols that provide services such as terminal emulation, the uploading and downloading of files, and access to pages published on the World Wide Web. Most implementations include applications that use the following protocols:

- HTTP (Hypertext Transfer Protocol)
- SMTP (Simple Mail Transfer Protocol)
- NNTP (Network News Transfer Protocol)
- FTP (File Transfer Protocol)
- Telnet

The Internet application layer protocols and the programs that run on them are discussed in more detail in Chapter 9, “The Widest Area Network: The Global Internet.”

Summary

Many protocols work in conjunction with the network/transport and higher-layer protocols to provide LAN and WAN communications. It takes many protocols, working together, to establish and maintain network communications between computers on a LAN or WAN. In this chapter, we discussed the three most popular network/transport protocol stacks: NetBEUI, IPX/SPX, and TCP/IP.

You learned that computer communication is based on binary numbers, yet humans prefer to use names to identify systems and resources. We discussed the available methods for resolving these “friendly” names to IP addresses that can be used by the computers.

This chapter focuses primarily on TCP/IP because it is the protocol of the global Internet and of most medium-sized to large LANs today. You learned about the suite of protocols that make up TCP/IP.

This chapter covers the basic concepts of IP addressing, and you learned about traditional address classes and CIDR, which is a new routing method that uses classless addressing. You learned about the two parts of an IP address, the role of TCP and UDP ports in getting a message to its destination, and how and why to subnet an IP network. You also learned about DHCP, which automatically assigns IP addresses to computers configured as DHCP clients.

The protocols used for network communications are complex. This chapter can provide only a limited overview. Many excellent books, Web sites, and classroom and online training courses cover these topics in depth.

Now that you have a basic understanding of how network protocols work, in the next chapter we discuss the largest and most complex network of all: the global Internet.

Further Reading

For an overview of IPv6, the next generation of the Internet Protocol, see playground.sun.com/pub/ipng/html/ipng-main.html.

An excellent resource on multicasting and the MBONE is available at www.savetz.com/mbone/ch3.html.

3Com provides a good basic overview of IP addressing, including subnetting, at www.3com.com/nsc/501302.html.

Review Questions

The following questions test your knowledge of the material covered in this chapter. Be sure to read each question carefully and select the *best* correct answer or answers.

- 1 Which of the following is the fastest and easiest to configure, but is a nonroutable, network/transport protocol?
 - a NWLink
 - b NetBIOS
 - c NetBEUI
 - d IPX/SPX
- 2 Which of the following parts of an IPX address is based on the MAC address of the device?
 - a Node number
 - b Network number
 - c Subnet number
 - d Host number
- 3 Which of the following are connection-oriented? (Select all that apply.)
 - a TCP
 - b UDP
 - c IP
 - d SPX
 - e IPX
- 4 Which of the following is true of IP addresses? (Select all that apply.)
 - a IP addresses are made up of two octets.
 - b An octet in an IP address that consists of eight bits.
 - c IP addresses are 32-bit binary numbers.
 - d IP addresses are usually notated in hexadecimal format.

- 5 How many host computers can a Class B network theoretically have?
 - a 16,384
 - b 254
 - c 65,535
 - d 2,097,152
 - e Over 16 million
- 6 What is the default subnet mask for a Class C network?
 - a 255.255.255.255
 - b 255.255.255.0
 - c 255.255.0.0
 - d 255.0.0.0
- 7 Assuming classful addressing, to what address class does the IP address 190.23.201.6 belong?
 - a Class A
 - b Class B
 - c Class C
 - d Class D
- 8 What is the routing method that uses classless addressing, with the network ID following the IP address indicated by a “slash x”?
 - a CIR
 - b CD-R
 - c CIDR
 - d RIP
- 9 What is the function of DHCP?
 - a It resolves NetBIOS names to IP addresses.
 - b It translates private IP addresses to public addresses.
 - c It resolves IP addresses to MAC addresses.
 - d It automatically assigns IP addresses to client computers.

- 10 Which of the following TCP/IP utilities can be used to view the cache of IP addresses that have been resolved to MAC addresses?
- a Ipconfig
 - b Iptrace
 - c ARP
 - d Nbtstat